Caligare Flow Inspector™ software
User Guide
Version 4.0

Applicability

This document applies to release of Caligare Flow Inspector Software, version 4.0.

Copyright Notice

© 2004-2006 Caligare s.r.o. All rights reserved. This software and the accompanying documentation are subject to copyright. You may not modify, adapt, translate, reverse engineer, decompile, or disassemble the software - or create derivative works based on it - without prior written consent of Caligare. Reproduction without permission prohibited.

Trademarks

Caligare Flow Inspector is a registered trademark of Caligare. All other trademarks are the property of their respective owners.

Disclaimer

Caligare does not give any guarantees or make any warranty or representation regarding this software and documentation, its correctness, accuracy, reliability, up-to-date, or otherwise. Neither Caligare nor anyone else who has been involved in the creation, production or delivery of this product shall be liable for any direct, indirect, consequential, or incidental damages (including loss of business profits, business interruption, loss of business information, and suchlike) arising from the use or inability to use the product.

Document Reference Information

Document Title: Caligare Flow Inspector ™ Software
User Guide
Version 4.0

Date: November 2006

Control: Caligare s.r.o. 21 Konselska
Prague 6, 18000 Czech Republic
E-mail: caligare@caligare.com

1 Introduction 3
What is NetFlow? 3
What is Caligare Flow Inspector? 4
Features and Benefits 4
Minimum System Requirements 5
Operating system 5
Hardware specifications 5
Minimum hardware requirements 6

2 Installation 7
Installation requirements 7
Installation in Debian distribution 7
Installation in RedHat and Fedora distributions 7
Installation in other Linux distributions 8
Installation script 8
Completing Setup 9
Debug Information 9

3 Getting Started 10

4 Configuration 12
Global settings 12
Device settings 14
Unit settings 15
Collector settings 16
Basic collector settings 16
Advanced collector settings 18
Anomalies settings 20
Anomalies - Collector settings 20
Anomalies - Global settings 22
Anomalies - Exclusions settings 23
Network settings 24
Application settings 24
Forwarding settings 25
Filtering settings 26
Image store 28
Host list 29
Port list 30
Country list 30
AS list 30
Group settings 31
User settings 32
Account settings 33

5 User Guide 35
Main screen - Overview 35
Data 36
History 36
Trends 36

Search 44
Interfaces 48
IP information 49
AS information 50
Graphs 51
Utilization maps 52
Profiles 55
Exports 55
Export list 55
Export status 56
Import list 56
Import status 56
Anomalies 56
Status 57
Engine 58
Devices 58
Units 58
Collectors 59
Last login 60
Tables 61
Database 62
Options 63
Help 63
Port database 63
License 64
Logout 64

6 Optimizing and tuning 65

Appendix 1: Configuring NetFlow Data Export 66
Configuring NDE on an IOS device 66
Configuring NDE on a CatOS device 67
Configuring NDE on a Native IOS device 68
Configuring NDE on a 4000 series switch 68
Configuring NDE on a Juniper router 68

Appendix 2: Frequently Asked Questions 70
Installation 70
Web interface 72
Other difficulties 76

Appendix 3: Network anomalies modules 77

Appendix 4: Third party software components 79
Apache web server 79
PHP 79
JPGraph library 79

1 Introduction

This document is a complete reference to the Caligare Flow Inspector (CFI) software, version 3.3. Its goal is to explain in detail the installation and configuration of the CFI software and illustrate different integration and application scenarios. CFI was created as a network monitoring and management solution, which collects NetFlow information from CISCO routers. This information is available for your review and/or analysis. This document is only a software manual and does not provide any assistance with any kind of devices/hardware itself. The document will be regularly updated. The latest version can be found and downloaded at: http://www.caligare.com/netflow/download.php If you have any questions about this documentation, please contact Caligare s.r.o.: caligare@caligare.com

What is NetFlow?

NetFlow is one direction only packet sequence between certain source and destination. Network devices (routers and switches) store and export all network data flows so they can be used for network management and network planning purposes.

NetFlow technology provides the data necessary to effectively analyze, trend and baseline application data as it passes through the network. It can then be exported to a reporting package and can provide the information necessary to manage critical business applications.

NetFlow records data consisting of information about source and destination addresses, along with the protocols and ports used in the end-to-end conversation. Caligare Flow Inspector uses this information to generate graphs and reports on traffic patterns and bandwidth utilization. NetFlow technology tracks the flow of IP packets as they enter the router through an interface. Each flow is unique and is identified by seven criteria; Source IP address, Destination IP address, Source Port number, Destination Port number, Layer 3 Protocol Type (TCP/UDP/ICMP/...), Type of Service (ToS), and Input logical interface, any variation in these criteria distinguishes one flow from another.

The types of information NetFlow can provide include:

1. *Network Monitoring in real time:* This technique is based on analysis of network packet exports, which are used for transparent display of dataflow going through the routers. This information then can be used for active detection and elimination of network problems.
2. *Application Monitoring and Profiling:* detailed statistics of used applications in different time intervals. Results from these statistics can be used for planning and specification of network topology. (For example: deployment and set up configuration of web server).
3. *User Monitoring and Profiling:* detailed statistics of individual network users. Statistics are used for effective planning and layout of load, deployment of cache servers, etc. It is also used for detection and solving potential security problems. User Monitoring and Profiling can tell you who the top users are,

how long they've been on the network, what Internet sites they've used, where on the network they go, what percentage of network traffic they use, what applications they use, and what are their usage patterns.

4. *Accounting/Billing:* Information about dataflow includes source and destination point information (IP address), number of transferred packets, bytes, time, used ports and type of service. This makes it suitable for detailed accounting among particular Internet service providers (ISP). ISP companies use these statistics for their services repayment, based mostly on the amount of data transferred.
5. *Network Planning and Analysis:* Network packet export can be used for network planning optimalization (e.g. who is communicating with who, planning and extension of backbone line and security rules). The main goal is to minimize the total price of network operations and maximize network performance, capacity and accessibility.
6. *Data Warehousing:* Network packet export can be archived for future analysis, making it possible to reconstruct all previous network traffic/activity. These services are very often used for statistics and graph generation by utilizing individual lines. It is also possible to estimate the services used by internal or external network users. This is especially valuable information for Internet service providers. Analysis of network packet export contains information about: what, where, with whom and how long they have communicated.

What is Caligare Flow Inspector?

Caligare Flow Inspector is a unique network software solution for companies, who need to plan, build, maintain and manage their network and at the same time keep their network more secure and efficient. **Caligare Flow Inspector is a web-based bandwidth monitoring tool that uses NetFlow data to provide detailed traffic statistics that help answer who, what, when, where of bandwidth usage.**

CFI software version 3.2.0 was engineered to create a secure network monitoring platform based on industry standards that will fit your existing security policies. The results are the ability to monitor in real-time, significantly reducing the time it takes to identify and troubleshoot. CFI keeps track of what is happening in your company ´s network, detecting attacks, and warning you of problematic network users. All information about network activities are archived in a central database.

Features and Benefits

Important facts:

- Having the ability to determine the true health of your network on a daily basis is a key component of your IT strategy and CFI gives you this wide visibility.
- Diagnose issues that degrade system performance, leading to quick resolution of issues without adding unnecessary infrastructure or bandwidth.
- Having the ability to access historical data, seeing patterns and trends, allows our staff to be more proactive in planning for the future.
- Having detailed information on where, by who and how specific applications are being used and how that usage affects the network.

- Using NetFlow data that is already present on company's routers and making real business decisions based on this information from a full enterprise perspective.

CFI provides you with:

- Detailed information about separate dataflow on the L3/L4 ISO/OSI network model.
- Hourly, daily, weekly and monthly statistics reports.
- The possibility of defining more statistics/characteristics according to user needs.
- Detailed and color graphs with tabs for every statistic.
- A definition of searching criteria in accordance to sub networks, used IP, used TCP/UDP port and detected application.
- A graph archiving possibility for future analysis.
- A definition of more users, where everyone can have their own settings.
- The ability to save search conditions in customizable profiles.
- Information about the status of devices and different ports through SNMP protocol.
- The ability to define descriptions of user applications.
- Convenient and proprietary monitoring of dataflow even on very large/extensive scale networks.

Minimum System Requirements

Operating system

CFI works under the all distribution of Linux (Debian, RedHat, Suse, Slackware, etc.), but preferred is Debian distribution. The Linux environment under which CFI software runs is considerable more stable and efficient, increasing the performance of the software.

Hardware specifications

It is very difficult to recommend optimal configuration, because good server performance depends on the amount of incoming data. Generally, there is an advantage in having adequate RAM memory and fast access to disc(s). The specification of your system depends on the number of routers sending network information to the CFI, as well as the level of actual router traffic.

Apart from the minimum hardware requirements set out below, is necessary to ensure that CFI should run on a dedicated PC or Server. The software is processor-intensive and in the case of very high loading (busy processor) it can cause problems in collecting NetFlow.

Manufacturer devices supporting CFI software are: Cisco Systems, Juniper, Extreme Networks and 3COM.

CFI supporting devices series (Cisco routers and/or switches): 1400, 1600, 1700, 2500/2600, 3600, 4500/4700, AS5300/5800, 7200/7500, Catalyst 4500, Catalyst 5000/6500/7600, ESR 10000, GSR 12000.

Please, ask your hardware supplier if your devices support NetFlow export.

Minimum hardware requirements

Following hardware requirements are the absolute minimum needed for the system to run:

- 256 MB RAM (RAM need to be increased to 1024 or more if you have a large network, or if more than one router is sending NetFlow traffic).
- 20 GB free hard disc space on the volume to which the database is installed, 100 MB free hard disc space on the volume to which the program is installed.
- Pentium III, 1 GHz or greater.
- Cisco router or any other that support NDE (NetFlow Data Export). The router and its IOS version must support NetFlow. For more information consult vendor's web pages.

These specifications will increase based on number of devices monitored. The highest computing performance is put on the database system. Computing requirements for the other CFI components are lower than for the database system.

2 Installation

Installation requirements

Apart from the Minimum System Requirements set out above, there are a number of things to check so as to ensure the best performance from Caligare Flow Inspector:

- Caligare Flow Inspector should run on a dedicated PC or Server. The software is processor-intensive and a busy processor can result in problems in collecting NetFlow data.
- We recommend the latest version of MySQL database server and client.
- Apache web server with PHP support.
- Installed PHP extensions php4-gd, php4-mysql and php4-snmp.
- System networking utilities: ping, traceroute, whois.

Before installation of NetFlow monitoring package, please, check if all required components are installed.

Installation in Debian distribution

Before installing stop any other or older NetFlow collectors! Installation in the debian environment is very easy. Download NetFlow package to directory /tmp and in shell type command:

```
dpkg -i /tmp/netflow_<version>.deb
```

where version is actual package version for example:

```
dpkg -i /tmp/netflow_3.2.0.deb
```

The Debian version runs installation script automatically. You can run this script later by typing nf_install in command shell. Continue to section Installation script.

Installation in RedHat and Fedora distributions

Before installing stop any other or older NetFlow collectors! Download NetFlow package to /tmp directory and in the shell type the following command:

```
rpm –i /tmp/netflow-.rpm
```

where version is actual package version for example:

```
rpm –i /tmp/netflow-3.2.0-1.rpm
```

After unpacking type nf_install in command shell to start configuration. Continue to section Installation script.

Installation in other Linux distributions

Before installing stop any other or older NetFlow collectors! In other Linux distributions, installation requires more manual input. Download NetFlow package to /tmp directory and in the shell type the following command:

```
tar -C/ -zxvf /tmp/netflow-<version>.tgz
```

where version is actual package version for example:

```
tar -C/ -zxvf /tmp/netflow-3.2.0-1.tgz
```

After unpacking type nf_install in command shell to start configuration. Continue to section Installation script.

Installation script

Before installing stop any other or older NetFlow collectors! From the menu you can select what part you want installed. In default all three parts are installed on same server.

- Press 1 to install database tables. This step is only used on the primary database server. Please enter license owner and license key. In case you want to use a trial version enter license key received by email. Trial version expires after 30 days. License owner and/or key can be changed via web interface. Now enter username and password for access to primary MySQL database. In the default installation of MySQL use username root and blank password. If the values are correct, the installation script will try to create a new database and all necessary tables. If you're doing an upgrade the old configuration tables will be backed-up.
- Press 2 to install web interface pages. Do this only on the web server. Now enter hostname of primary database. In default don't enter any value, because the primary database is on the same machine as the web server. Next parameter to enter is the database port number (the default value is empty). Next enter the username and password, use the same username and password as are configured in the primary database. Please, refer to MySQL documentation (http://www.mysql.org/doc) to view how to create users or change passwords in MySQL database. The script will now create a new configuration file for the web part of NetFlow monitoring software and try to find the apache configuration file. If successful, the script will include a web part into the apache configuration and then restart the web server. If unsuccessful, you must include the file /etc/netflow/apache.conf to your web server configuration manually.
- Press 3 to install the collector. In case you want to use more collector servers, repeat this step for all of them. During the collector installation part, after entering database parameters, you will see a list of configured units. Each unit is a corresponding server on which you can run one or more collectors. Enter the unit ID on the installation computer. This unit ID is unique and can be used by only one server! In other words, each server has unique unit ID. If you want

to use more servers as collectors, you must enable MySQL networking option (see MySQL documentation or Appendix 2 how to enable networking) before creating new units via the web browser.

- Press 4 to finish.

Run NetFlow collector process via command:

```
/etc/init.d/nfcd start
```

on all servers whose collectors can run. If nfcd process isn't running see syslog for error messages or troubleshooting section.

Completing Setup

When setup is complete, launch web browser and open address *http://your_webserver/netflow* to verify that the system is running.

To login use default username: **admin** and password: **nfadmin**
We recommend changing administrator password as soon as possible.

You can now proceed to configuring the system. The Getting Started section of this manual covers the essentials of getting NetFlow monitoring software up and running.

Debug Information

Debug information helps us determine where the problem was with your un-successful installation. Log into Linux system console and run the following command:

```
nf_debug
```

This command creates a debug file, which will be sent to our support email address. You can display this file via software web interface (menu Help->Debug file).

Debug file contains:

1. MySQL configuration - all important tables are dumped.
2. Configuration netflow files.
3. IP address setup, default gateway, etc.
4. Time used in the system with time zone information.
5. Up and running processes.
6. Incoming packets dump (tcpdump).
7. List of opened network connections (netstat).
8. Report from the system log file.
9. MySQL library version, PHP and web server configuration etc.

3 Getting Started

Installation and configuration of Caligare Flow Inspector is simple. This section addresses the few essential steps required to collect and display the NetFlow information from your network. More detail for each step is available in subsequent sections of this manual.

1. Set up NetFlow Data Export (NDE) on your router(s) or L3/L4 switch(es). Appendix 1 gives a quick guide on setting up NetFlow Data Export on Cisco devices.
    - For more information on this, refer to your router documentation, or go to http://www.cisco.com/go/netflow.
    - Set the destination of the NetFlow traffic to the IP address of the NetFlow collector workstation.
2. Install NetFlow monitoring software on the workstation as shown in the previous section.
3. You can access the web-based interface of Caligare Flow Inspector using a web browser. For access to web interface use the following address: *http://<your_webserver>/netflow*, <your_webserver> is the IP address or hostname of the web server where NetFlow web part is installed.
4. Log into system using username admin and password nfadmin and select Options menu.
5. Most of the configuration defaults will allow you to start collecting data, but there are some items that require setup:
    - Device settings: When a router or switch sends NetFlow Data Exports to the monitor it is important to setup the IP address and SNMP community string for resolving interface names. This step is necessary in case, the devices send data to the collectors, which are all listening on the same port. It's recommended using read-only SNMP community for security reasons.
    - Collector settings: Add new collector. It's recommended using standalone collector for each router. When creating a new collector select the unit on which you want it to run, listening port (e.g. 2000), number of hourly tables which will be stored (e.g. 32), number of daily tables (e.g. 31), number of weekly tables (e.g. 4) and number of monthly tables (e.g. 3). Don't forget to enable collector. Advanced parameters can't be modified later! You can choose aggregation steps and which items you want to store in the hourly, daily, weekly or monthly tables. You can select the format of stored data and which categories you want to store (e.g. source IP, destination IP). Properly formatting the tables helps save disk space, because you can limit which items are able to be viewed and stored.
    - License settings: If you have received a full license with this product, it should be loaded via menu "Help->License". This product can't run without license key. See web pages download section for getting trial demo key.
6. Minimum configuration is now complete. For more configuration information, read the Configuration Guide. Graphs will be available in seconds after

starting Caligare Flow Inspector software. After successfully login click on menu "Options" for configuration or menu "Data->Trends" for view graphs.

4 Configuration

Any configuration options are done in Options menu. Visibility of options depends on your access rights, so that the common user can't see many of these settings. Latest released version has 10 option submenus:

- Global - you can specify administrator email address, default skin, etc.
- Devices - manage NDE devices (routers and switches).
- Units - manage servers on which you run NetFlow collectors.
- Collectors - manage all collectors, listening ports, number of stored tables.
- Anomalies - configure network anomalies detection.
- Networks - define your network or foreign IP networks.
- Applications - define rules for application recognizer.
- Forwarding - define rules for forwarding NDE to other destinations.
- Filtering - define rules for forwarding NDE to other destinations.
- Image store - upload and manage images for using in the utilization maps.
- Port list - manage port name database.
- Country list - manage database of countries.
- AS list - manage autonomous systems database.
- Groups - manage groups of users and their access rights.
- Users - manage users, sets graph resolutions, skins, etc.
- Account - change account values of actual logged user.

Global settings

In the "Global settings" you can change the skin of the web interface, default SNMP community string, email setting, LDAP authentication and anonymous login extension. In the global options you can enable or disable checking for new versions of the software by displaying last the logins and/or displaying license(s) expiration warnings. If you select the option "email logins to administrator", all users who login will be reported to the administrator's email address.

If you select the option "Display overview statistic", you will see how many bytes, packets and flows each collector parsed. You can also enable or disable generation of utilization graphs via the option "Display overview utilization graph". The option "Utilization graph history" gives you the ability to determine how long the history will be displayed. This value can be set between 30 and 360 minutes.

You can import graph icon image(s) and assign your company logo to all graphs. See "Image store" menu for more information about upload images. Parameters Position-X and Position-Y indicate the position where a graph icon will be displayed. The position can be specified as either absolute coordinates or as a fraction of the width and height respectively. A negative value means that the anchor will be right or below the icon. E.g. Position-X=-1 and Position-Y=-1 mean that your graph icon will be displayed in the bottom right corner.

We recommend enabling DNS caching option. If you enable DNS caching, all domain name resolution queries will be cached and stored on your system disc. Positive and

negative timeout parameters give you the ability to set how long queries will be stored in the cache.

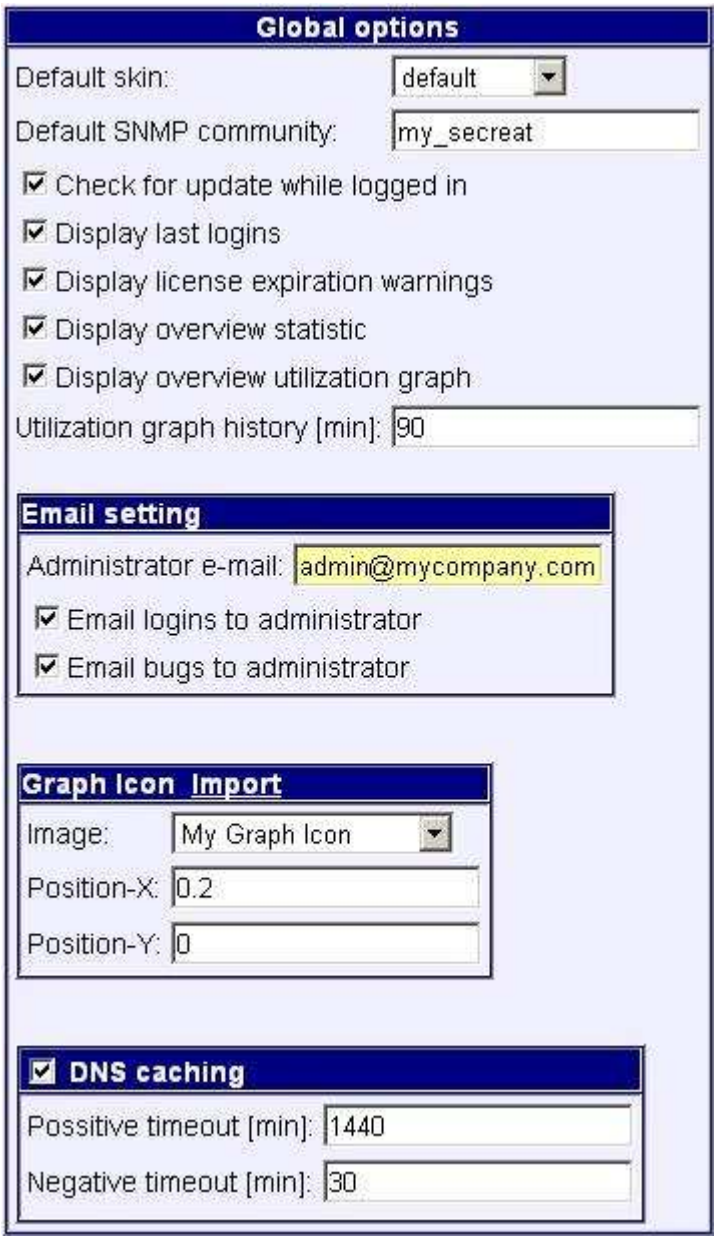

Figure: Global settings window.

Version 3.2.2 implemented LDAP authentication extension that uses LDAP server for user authentication. For example you can use the following LDAP server URL: ldap://ldap1.mycompany.com and LDAP bind DN: uid=%u,ou=people,dc=mycompany,dc=com A percent sign (%), followed by character (u) is replaced by username.

Version 3.2.4 implemented an external authentication extension that uses local system scripts or programs for user authentication. The program or script reads the entered password on a standard input, and if the user is authenticated the return code is sent back as zero. A non-zero return code means that the user entered a bad password or script error. For example you can use the following command:

```
/usr/bin/nf_auth_pam %u
```

A percent sign (%), followed by character (u) is replaced by username. Program nf_auth_pam uses LINUX system authentication module (PAM). Netflow monitoring package also includes script nf_auth_smb with which you can authenticate users via your windows domain controller. For more information about windows authentication see /usr/bin/nf_auth_smb file.

If you want to enable anonymous login, create new account that will be used for anonymous login first. In the global settings enable anonymous login and assign an anonymous username to anonymous account.

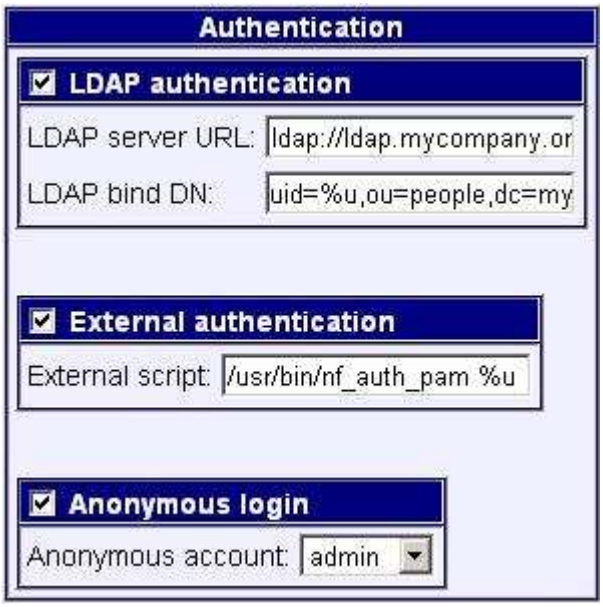

Figure: Global settings window - Authentication settings.

Device settings

In the device setting you can manage all NDE devices, such as routers or L3/L4 switches. If you want to see the state of various interfaces and/or interfaces names, it is necessary to set up the SNMP parameters as a community string and the IP address of the device. We recommend using a read-only community string for security reasons. The IP address is the same as that used for NetFlow data exports. In most cases use the IP address of the interface closest to the NetFlow collector. The Appendix 2 section will show you how to find this IP address.

In device setting you can modify sampling values. If you're using NetFlow sampling on the router, every N packet is added to the info flows, so in total sum you see only ~ 1/N data rate. When using this option all incoming traffic will be multiplied by this constant. You can also resample flows in the collector, which helps when the database is overloaded. You can set resampling to level 5, so that every fifth flow will be counted and the remaining four will be discarded.

Option "Automatically creates new collectors" which causes that master process listen to all incoming packets. If the source IP address is the same as the IP address

of configured device, this option will automatically create a new collector, that listens to this traffic. If this option is available, we recommend creating all the collectors manually.

In the list of devices you can use the "Interfaces" command. This command displays a new window that allows you to enter a name and comment for that particular interface.

In the list of devices you can use the "config file" command. This command creates a netflow configuration for the selected device. Netflow configuration generation is supported only for IOS/CatOS Cisco compatible devices and those which are accessible via SNMP protocol.

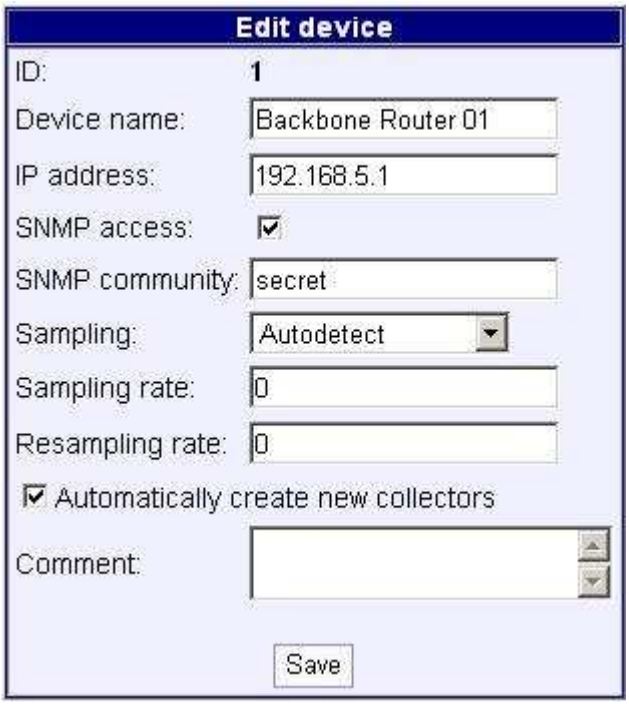

Figure: Device settings window.

Unit settings

If you are using the all-in-one server, you don't have to create a new unit, because the first unit is already predefined. If you want to use more servers with the collectors, you first need to create new units, one unit for one server. The unit identification number (unit ID) is very important. This number must correspond with "unit_id" value in the configuration file of the NetFlow collector (*/etc/netflow/nfcd.conf*).

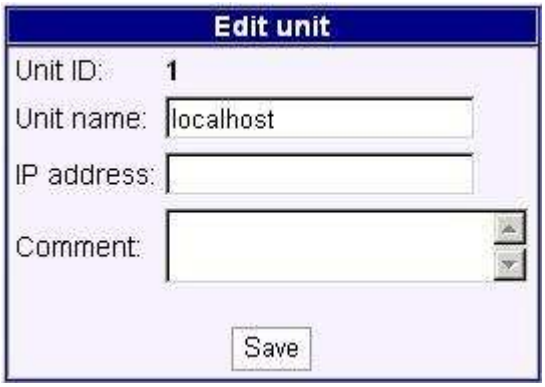

Figure: Unit settings window.

Collector settings

Basic collector settings

Collector settings are the most important option. For each collector you must set up the listening port, number of tables that will be created and stored and the associated NDE device(s). This has to be set up on the unit (server) that you want run on the selected collector. Listening port will begin in an interval between 1024 and 65535. Commonly the used value for the listening port is 2000 and must correspond with a value configured on the NDE device. The number of tables depends on your disc space and incoming data flow. For example routers with ten 100Mbs interfaces and a 20GB disc, the optimal values: for hourly tables is 48, for daily tables 31, for weekly tables 4 and for monthly tables 3.

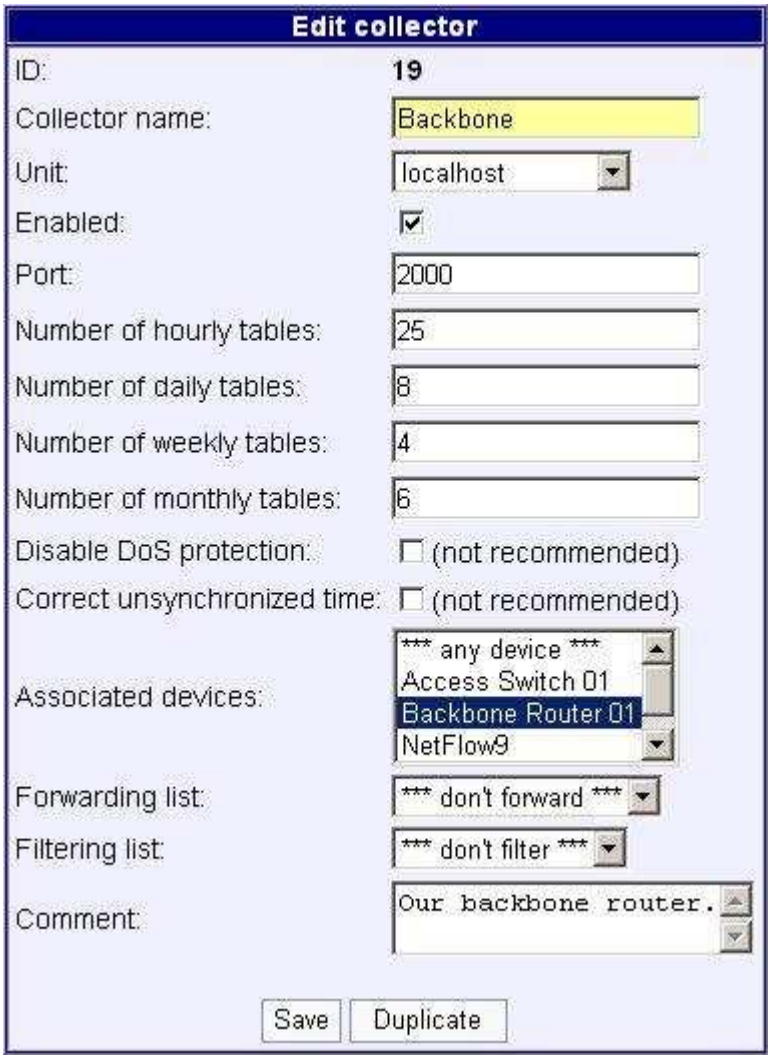

Figure: Basic collector settings window.

Collector has denial of service (DoS) protection. If the collector detects a big increase in traffic, it will stop all short flows for 1 minute. If that does not help, it will block all flows. This feature protects your database from overloading. You can disable this feature if your traffic is too shaky.

If the time between collector server and exporting device is unsynchronized, flows that contain the wrong time will be dropped (see the menu Status->Collectors and 'Dropped flows due to corrupted time' counters). You can correct the wrong time by changing the collector settings (option correct unsynchronized time). In most cases the source of the problem is a different/wrong time zone setting or wrong time set up on exporting device. The collector by itself analyzes each flow and if there is a difference between the flow time and the collector's time by more than 12 hours, the flow time is replaced by the collector's time.

It's possible to configure a forwarding list if you want to forward NDE to other destination(s). Before enabling the forward or filter feature, the forward or filter list must be defined via the "Forwarding settings" or "Filtering settings" menu. In case you want to resolve interface names it is important to associate a NDE device with the collector. Don't forget to enable the collector. Advanced settings can be changed only when you define a new collector!

Advanced collector settings

In Advance collector settings you can select the short aggregation step. For hourly tables this step can't be set up (it's always one minute). For daily tables it can be one hour (default) or 30 minutes or 10 minutes. For weekly tables it can be one day (default) or 12 hours or 6 hours. For monthly tables the only possible values are one day (default) or 12 hours.

Reduce factor

Automatic size reduction is used in the Netflow software. This means that uninteresting (low volume) flows are not inserted into the aggregated tables (daily/weekly/monthly). The reduce factor parameter gives you the flexibility to set the amount of traffic that will be dropped. For the aggregation from the hourly tables into the daily tables there is a maximum of 3% total volume dropped, for the aggregation from daily to weekly or monthly tables there is a maximum of 1% dropped.

There are several exceptions to the rule.

1. If the number of rows in the source table is less than 200,000, then no size reduction is used.
2. If the number of aggregated rows is less than 5% rows, in the source table, no size reduction is used.
3. Aggregated table must have flows that are higher than 200kB.

A reduction factor value is set as a percentage (from 0.0 to 20.0). A zero or empty value means that the system will use the default settings. You can disable the size reduction feature by setting this value to “–1”. If you disable size reduction you risk that the collector will create huge tables whereas queries may fail and the overall system may become unstable!

Setting the correct format of the tables can be very useful. For ISPs, the BGP AS numbers and next hop address can contain some interesting information. Another interesting feature is setting up accounting of source and destination interfaces on a backbone router. This setting will give you freedom to choose what you want to monitor. The more items that are selected can dramatically raise the amount of space required to store these records. Daily tables depend on hourly tables, so the format of the daily tables can be the same or reduced in format compared to the hourly table. Weekly and monthly tables depend on daily tables.

You can use one of predefined formats:

- Basic (basic fields as IP addresses, protocol, ports, interfaces, application and next hop)
- Hybrid mode (same as Basic, but it adds exporter IP address – useful for devices that works in the hybrid mode)
- BGP (same as Basic, but it adds autonomous system information)
- Security (same as Basic, but it adds TCP flags and type of service fields into hourly tables)

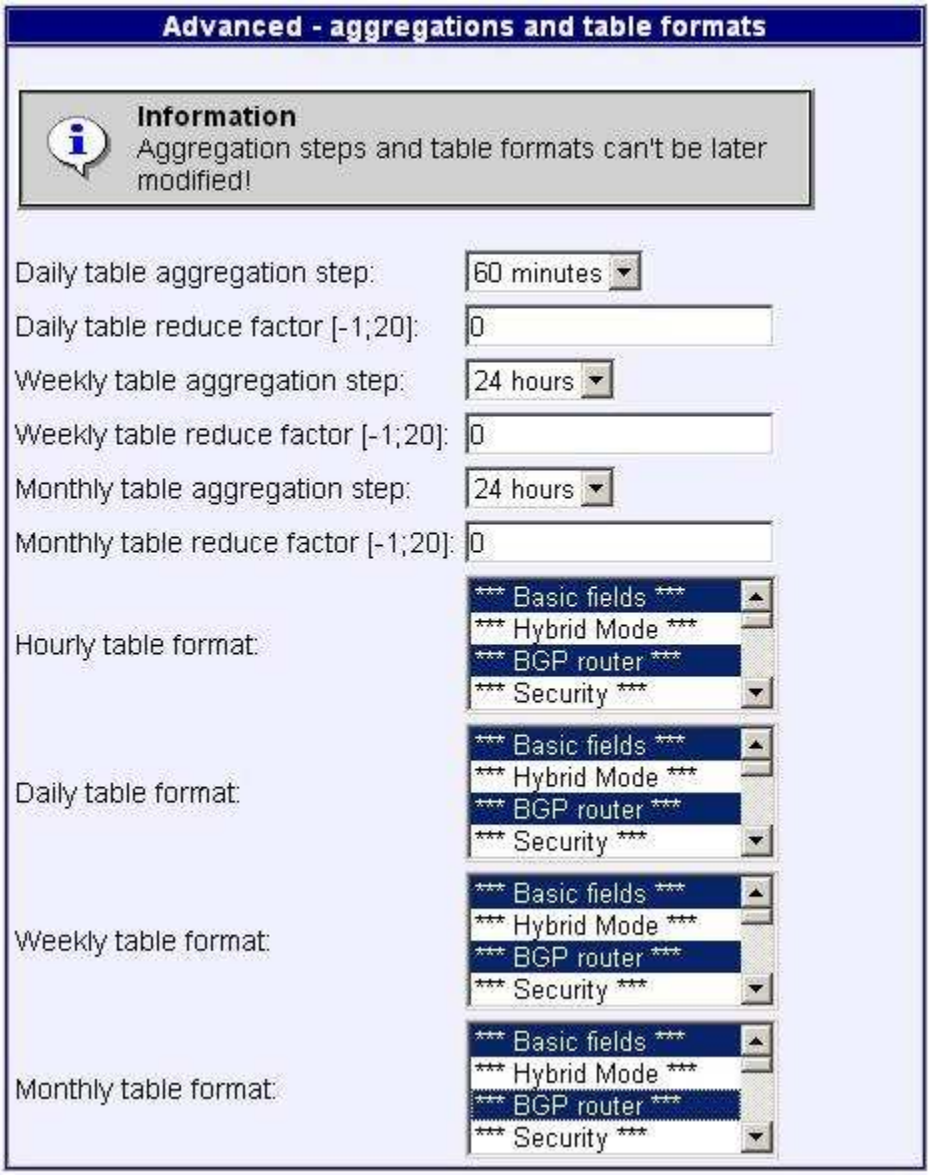

Figure: Advanced collector settings window.

Default format of hourly table is:

- Source IP address
- Destination IP address
- Application
- Protocol
- Source port
- Destination port
- Source interface
- Destination interface
- Next hop IP address

Default format of daily table is:

- Source IP address
- Destination IP address

- Application
- Protocol
- Source port
- Destination port
- Source interface
- Destination interface

Default format of weekly table is:

- Source IP address
- Destination IP address
- Application

Default format of monthly table is:

- Source IP address
- Destination IP address

Anomalies settings

Packet sniffer is more a troubleshooting tool than a specific tool for constant netflow monitoring. Packet sniffer allows you to capture every packet and store it on your hard disk. Let's say you want to do 24 hour monitoring - 7 days a week, this way you need an incredible big hard disk.

Netflow monitoring collects statistics not the whole packet, which is why this method is more suitable for constant monitoring.

This current software version supports base network anomaly detection such as network and host port scanning, ICMP and TCP/SYN flooding detections, and detection of network games and peer-2-peer applications. Most of the modules use heuristic detection methods - for every anomaly there is a specified probability of incident.

Anomalies - Collector settings

If you want to run network anomalies (NA) detection it's required that you enable the NA for every collector. NA detection consumes a lot of CPU and memory so be careful when enabling this option.

This software also enables you to specify internal network IP address ranges for every collector. If NA module detects that incident is related to the internal network it gives the anomaly higher severity. IP address range can be specified in the following formats: single IP address (10.1.1.1), domain name (web.mydomain.com), list of IP addresses (10.1.1.1, 10.2.1.1, web.mydomain.com), range of IP addresses (10.3.1.1-10.3.255.255), IP networks (10.0.0.0/8, 192.168.0.0/16), exclude range of network (10.0.0.0/8, !10.1.0.0-10.5.255.255). The list of IP addresses has to be separated by a comma.

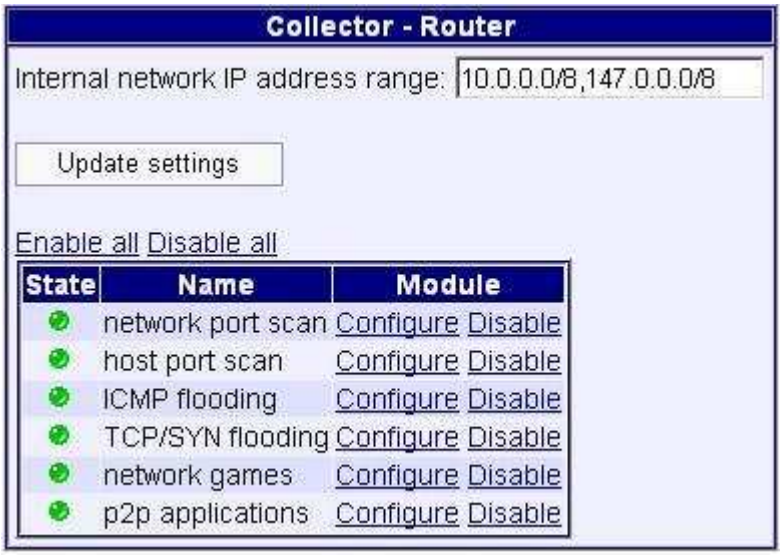
Collector - Router
Internal network IP address range: 10.0.0.0/8,147.0.0.0/8
Update settings
Enable all Disable all
State Name Module
network port scan Configure Disable
host port scan Configure Disable
ICMP flooding Configure Disable
TCP/SYN flooding Configure Disable
network games Configure Disable
p2p applications Configure Disable

Figure: Anomalies - Collector settings window.

You can configure severity of anomaly for each network module. Severity is specified as a function of probability and the number of anomaly occurences. For example you configure 10 occurrences for important severity. Analyzing software may assign important severity if it detects more than 10 occurrences with 50% probability or 5 occurrences with 99% probability or 20 occurrences with 1% probability. Occurrence value '-1' means that you don't want to generate a severity for this anomaly.

Other settings are module dependant (for example: sensitivity, minimal number of observed destinations, used TCP/UDP ports etc).

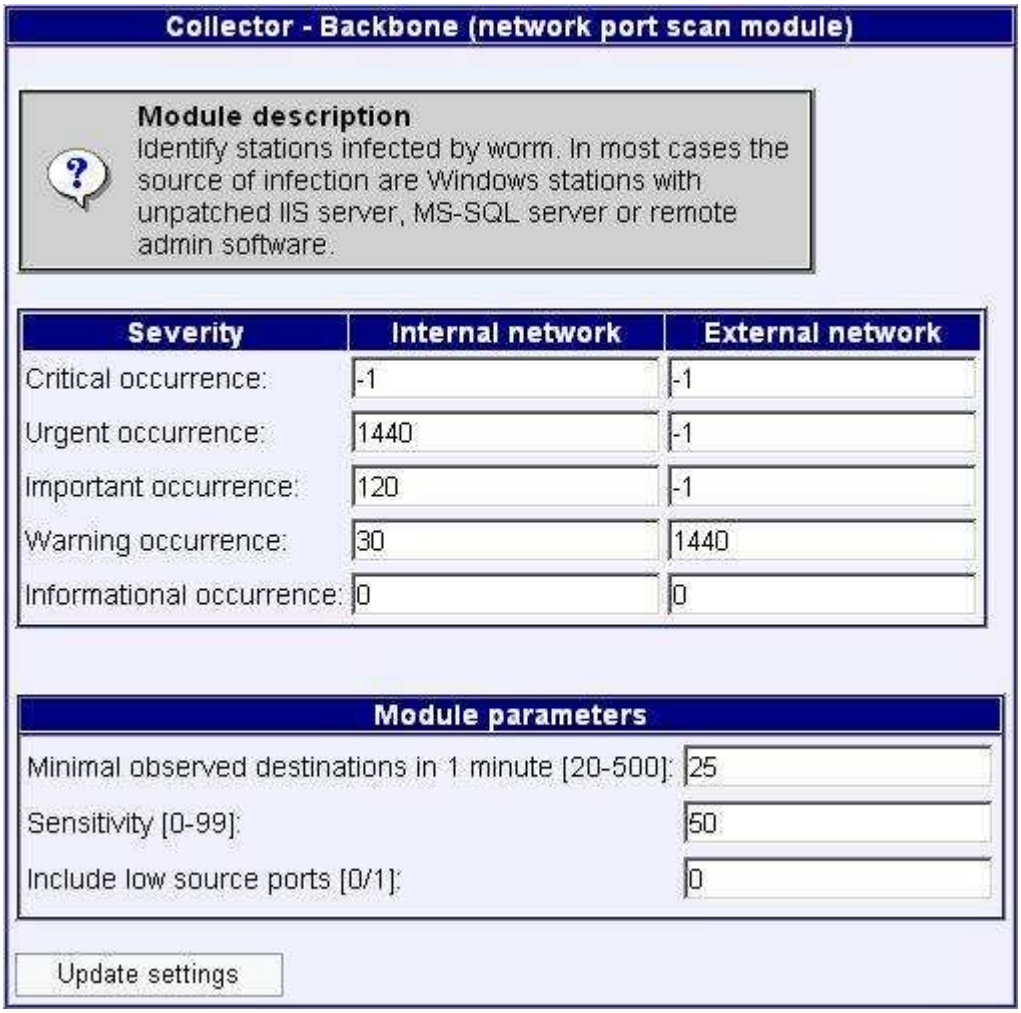

Figure: Anomalies - Module settings window.

Anomalies - Global settings

In the 'Anomalies global settings' you will be able to change the report parameters, intervals for removing old incidents and incidents colors. By clicking on any incident you can select its reporting by email. Incidents can be reported to two email addresses, the first one is for internal network incidents and the second one is for external network incidents.

You may also specify email subject (i.e. Network incident %INC – directive %INC is replaced by incident number), email body, header and tail. Maximum size of email is 256 characters (including incident detail text).

The incident removal option allows you to choose interval for old incidents removal. The first option 'clean new incidents' specifies interval for the new incidents removal. 'Clean new incidents' value is in interval 1-91 days. The second option 'clean other incidents' specifies the interval for removing any other incident state instead of 'archive' state. 'Clean other incidents' value is in the range of 1-200 days.

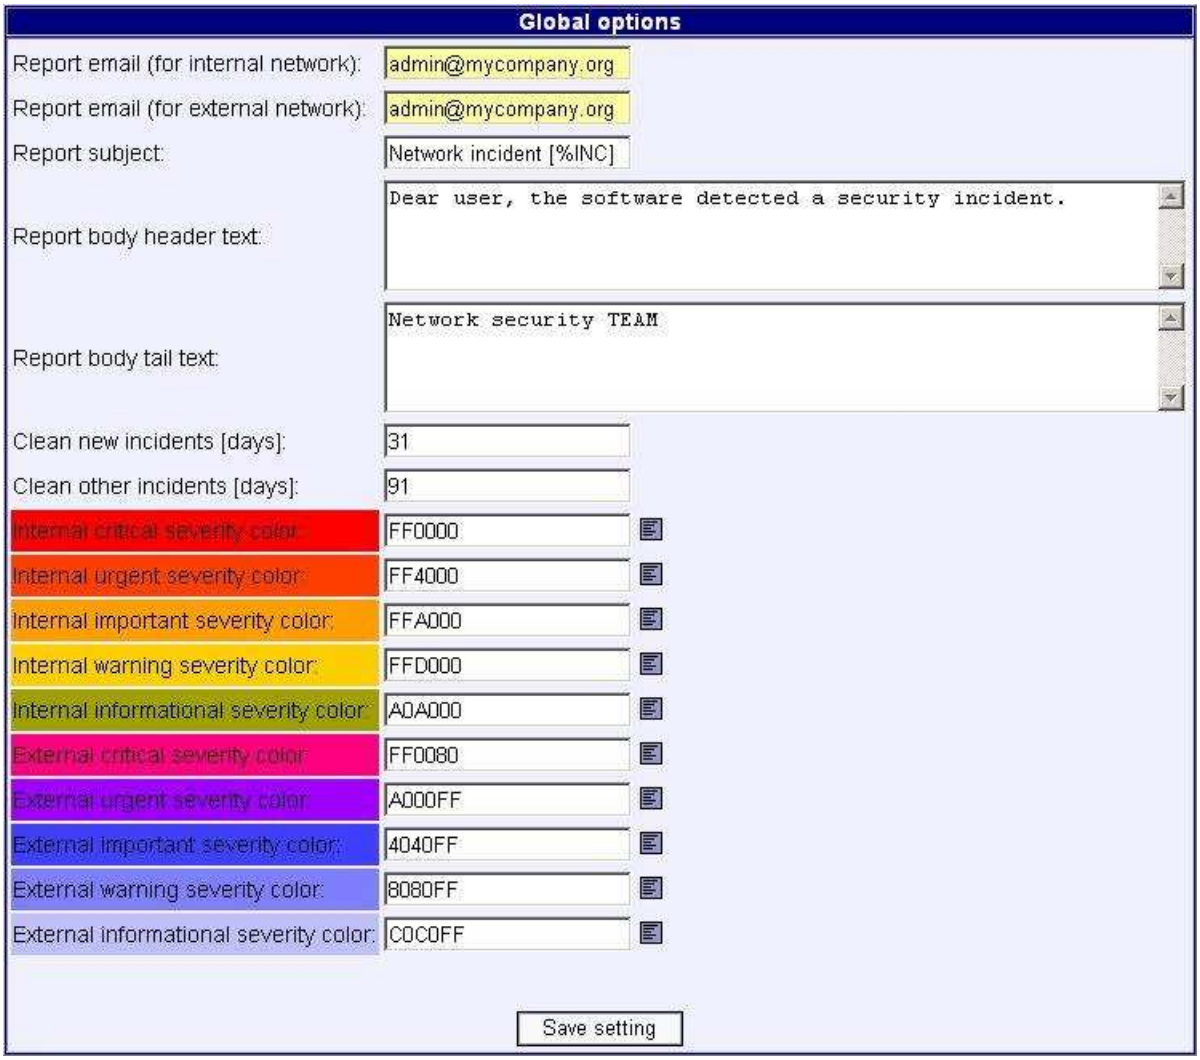

Figure: Anomalies - Global settings window.

The next, very helpful, feature for incident marking allows you to choose the incident colors. Network anomaly detection software uses 5 severities: critical, urgent, important, warning and informational. You may select a color for any severity. Color is defined as six hex digits (RGB format so called red-green-blue format). Some examples of color codes: red - FF0000, green - 00FF00, blue - 0000FF, cyan - 00FFFF, magenta - FF00FF, yellow - FFFF00 etc.

Anomalies - Exclusions settings

'Exclusions' screen shows you a list of network anomalies exclusions. If you want to exclude some anomaly, click on the source or destination of network anomaly in the main menu 'Anomalies' and select the exclude action. Exclusion can be active for 24 hours, 3 days, 7 days, 31 days and forever. You can also select for which network module you want activate exclusion etc.

Be careful when adding a new exclusion(s), too many exclusions may heavily load the system!

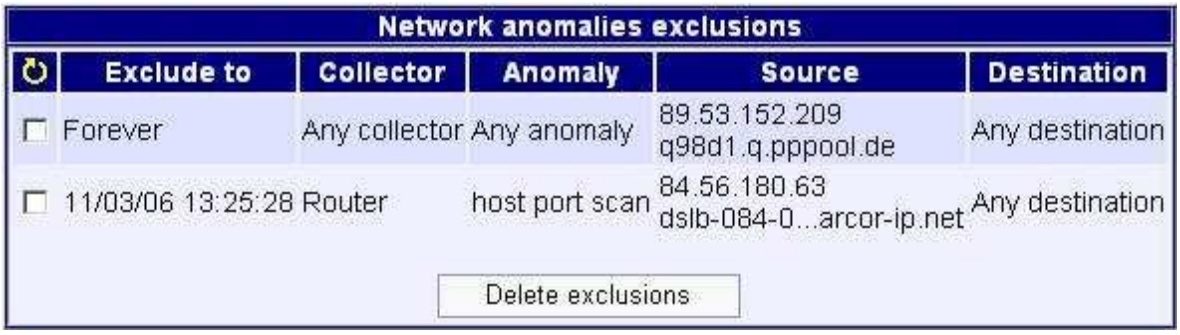

Figure: Anomalies - Exclusions settings window.

Network settings

The main purpose of this menu is to define IP ranges and name them. Defined networks can be used in menu "Data->Trends" and "Data->Search".

Some examples:

- Single IP address (10.1.1.1).
- Domain name (web.mydomain.com).
- List of IP addresses (10.1.1.1, 10.2.1.1, web.mydomain.com).
- Range of IP addresses (10.3.1.1-10.3.255.255).
- IP networks (10.0.0.0/8, 192.168.0.0/16).
- Exclude range of network (10.0.0.0/8, !10.1.0.0-10.5.255.255)

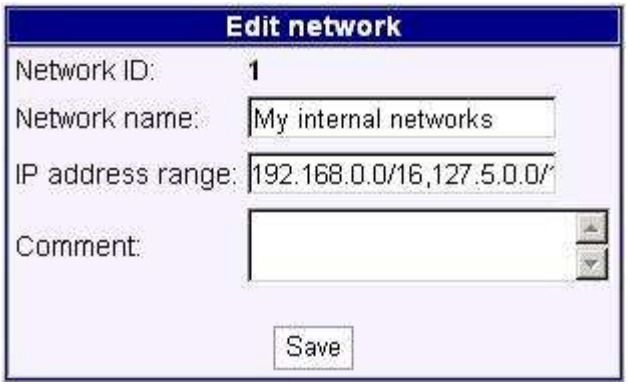

Figure: Network settings window.

All previous types can be combined. Field separator can be a comma or semicolon. You can also use an exclamation character '!'. This character excludes single IP or a range of IPs from the list. When you use IP address ranges, domain names can't be used!

Application settings

Caligare Flow Inspector contains a special application detection module (ADM). The ADM detects dynamically assigned ports.

Figure: Application settings window.

You can define your own application via the applications settings menu. One of your applications may contain more application rules (see picture bellow). The ADM uses system file */etc/services* to detect non-specified applications, but in this file you may specify only a single UDP or TCP port with the application name. The ADM module is very time-consuming, so be careful when you define more rules.

The ADM module can store a detected application into the field "app". In the raw data you can see "app" field values in these intervals:

- 0-65535: TCP ports in range 0-65535. Number corresponds to TCP port number.
- 90000-90255: not TCP, UDP or ICMP protocol, value 90047 means that in flow is used for protocol 47 (GRE).
- 99999: Source IP address is same as destination IP address or source and destination ports are zero. They are possible spoofed IP addresses and unknown application.
- 100000-165535: UDP ports in range 0-65535, a 100189 value means that the UDP protocol and 189 is the corresponding port number.
- 200000-265535: Used for ICMP protocol.
- 300000-unlimited: Used for applications defined via application settings.

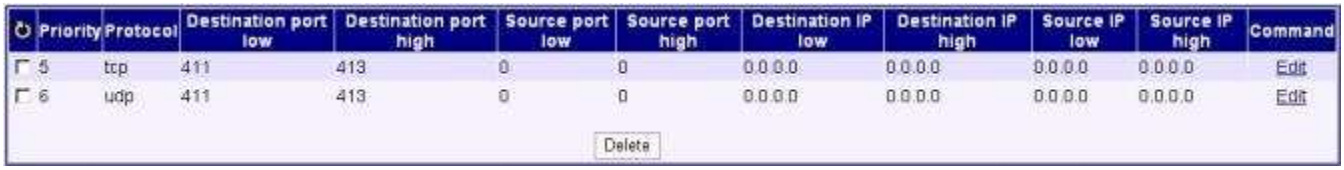

Figure: Application rules window.

Each rule contains priority, protocol (UDP or TCP). Other fields contain the destination port range, source port range, destination IP address range and source IP address range. You can fill up only some of these fields, the others are remain unfilled or have a zero value (it mean match any). In the example above, there are two rules, one is for the UDP and the other one is for the TCP along with a destination port (which has a range from 411 to 413), all other fields are zero. (it mean match any). The application used for example above is direct connect.

Forwarding settings

In the forwarding settings you can specify a list of destinations where you can forward NDE.

Figure: Forwarding settings window.

The setting is very similar to the application settings, with one difference. In the rules editor you can specify destination IP address and destination port. The picture below shows NetFlow traffic that will be forwarded to IP address 10.1.1.20 and port 2000. Version 3.3.0 implements source IP address spoofing. If you enable this feature, the collector modifies the source IP address of forwarded packets to the IP address from which the packets were originally received. This feature cannot be used where Cisco reverse path check feature is enabled. Collector settings assign the created forward list to the collector that will forward the NDE.

Figure: Forwarding rules window.

Filtering settings

Version 3.3.0 implements a flow-filtering feature that uses certain rules/conditions. In each rule you can specify conditions and actions that are to be performed when conditions match a certain flow. You can use the following types of actions; deny, modify or allow. The action "deny" drops flow. Dropped flow is not stored into the database. "Deny" action can be used for removing unwanted traffic from accounting. The "modify" action replaces flow with values that are specified in the set fields and continues with flow filtering. "Allow" action works similarly to the "modify" action, but it doesn't continue to filter flow. In other words allowed flow is stored into database, modified flow may or may not need to be stored into the database (it depends on which allow or deny rules follow). The default rule is to permit any flow.

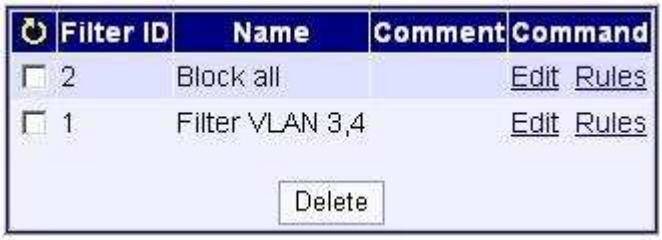

Figure: Filtering settings window.

For each rule you can specify up to 10 conditions and 10 "set fields". There are no rule limits, but be very careful in how many rules and conditions you create. Filtering consumes a lot of CPU time!

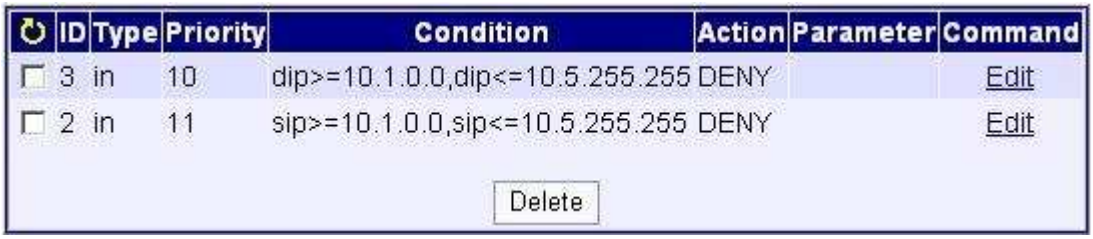

Figure: Filtering rules window.

Filtering feature can also be used for replacing a source IP address. If you are receiving netflow traffic through a netflow forwarder, incoming netflow shows IP of the forwarding device, instead of the IP address of the router that sent this information. The filtering feature has the ability to change the IP address which will correspond with the original device that sent the information. Collector settings assign the created filter list to the collector that will filter the NDE.

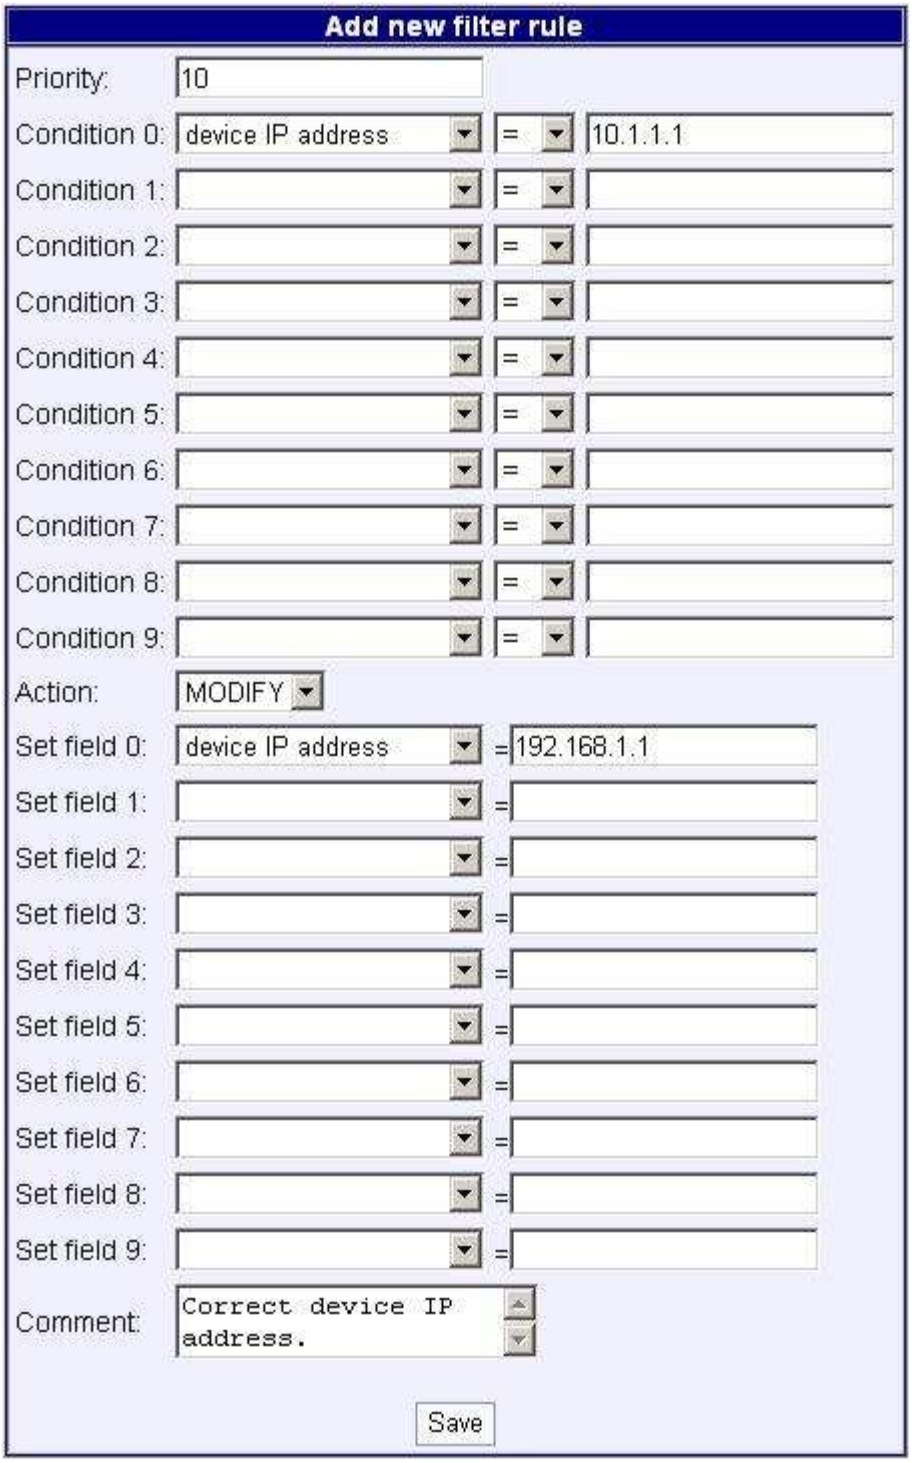

Figure: Filtering rules - condition window.

Image store

In the menu “Image store” you can manage and upload images. Uploaded images can be used in the utilization maps. Size of uploaded image is only limited by PHP and MySQL settings. If you want to use a big image (>8MB), modify the PHP options: *post_max_size*, *memory_limit* and *upload_max_filesize*. The maximum supported image size is 16MB. Uploaded images are base64-encoded and stored in the MySQL database. This encoding is designed to make binary data survive transport through

transport layers that are not 8-bit clean. Base64-encoded data takes about 33% more space than the original data.

Before storing image into a graphic (GD) library database, check if the graphic format is supported. Supported image formats are JPEG-JFIF Compliant format [JPEG], CompuServe Graphic Interchange format [GIF] and Portable Network Graphics format [PNG].

For each image you can specify name, group and type. Only three groups are recognized for use in the utilization maps. First is the "_UTILIZATION_MAP" (used for background image), the second one is the "_UTILIZATION_OBJ" (used as object image) and the third is the "_GRAPH_ICON" (used for graph icon). Utilization maps can work with transparent colors. Magenta color (#FF00FF in RGB model) is used as the transparent color.

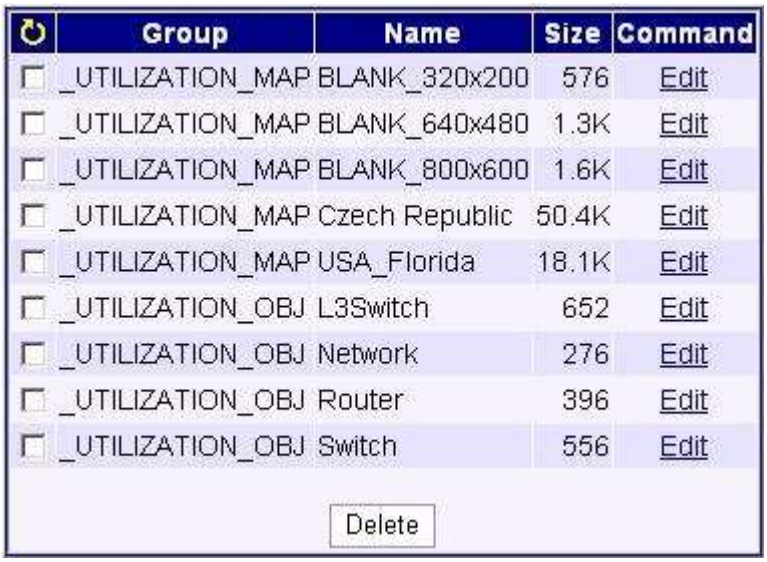

Figure: Image store.

Host list

The "Host list" feature enables you to define a certain name for any IP address. This host name assignment will be later used in the "Trends" or "Search" menu.

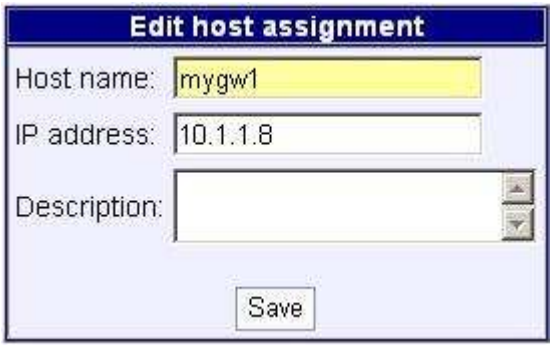

Figure: List of hosts - dialog window.

Port list

The "Port list" feature enables you to define a certain name for any port number. This port number assignment will be later used in the "Trends" or "Search" menu. The port name is converted in to the lowercase.

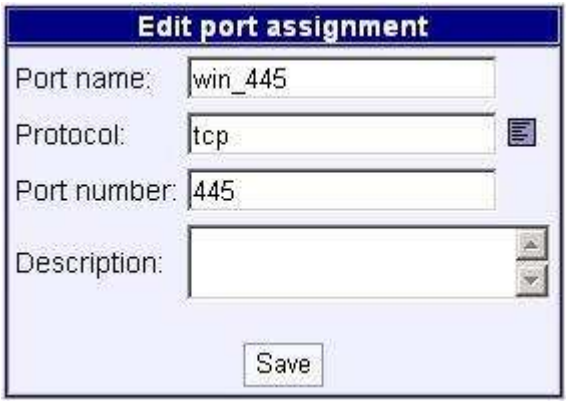

Figure: List of ports - dialog window.

Country list

The "Country list" option enables you create a new country name and assign IP address range to this country. The software has 233 countries internally stored and many IP address mapping. This setting overrides the internal country database.

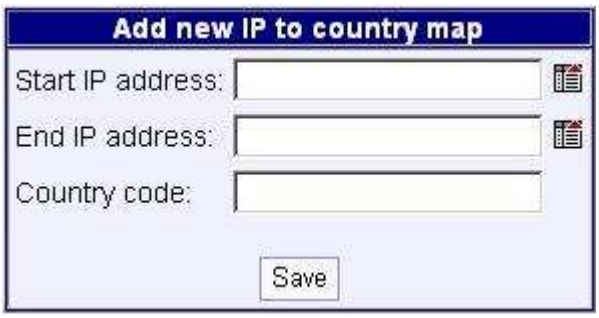

Figure: List of countries - dialog window.

AS list

The "AS list" is used for creating a new autonomous system number and assigns an IP address range to this autonomous system. This setting overrides the internal autonomous system database.

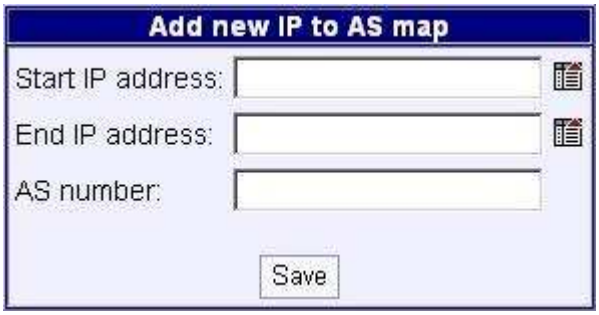

Figure: List of autonomous systems - dialog window.

Group settings

Main purpose of this menu is to create a named group of users and to assign rights to this group. Available rights are:

- Administrator - you have all rights. Only user with administrator rights can create new groups and users.
- Configuration - this enables you to edit all submenus in the Options menu.
- Collector maintenance - this enables you to edit collector settings.
- View status - enables access to menu "Status" and view status of the collectors and database.
- Search statistics - enables run "Data->Trends" and "Data->Search" statistics.
- Export data - enable export data from Trends and Search statistics.
- Profiles - enables you to save search profiles.
- Shell commands - enables you to run shell commands from the web interface as ping, traceroute and whois to get information about IP addresses or autonomous systems.
- Utilization maps - enables you to create a new utilization maps, objects and paths.

You can set the traffic view restrictions for each user group. If you assign restriction rule to a user group, only the collectors or data matching condition(s) will be displayed. In the "add group restriction rule" you can specify the restriction type, conditions and on which of collector you want to apply this rule. Format of the condition field depends on restriction type (see caption "Trends conditions").

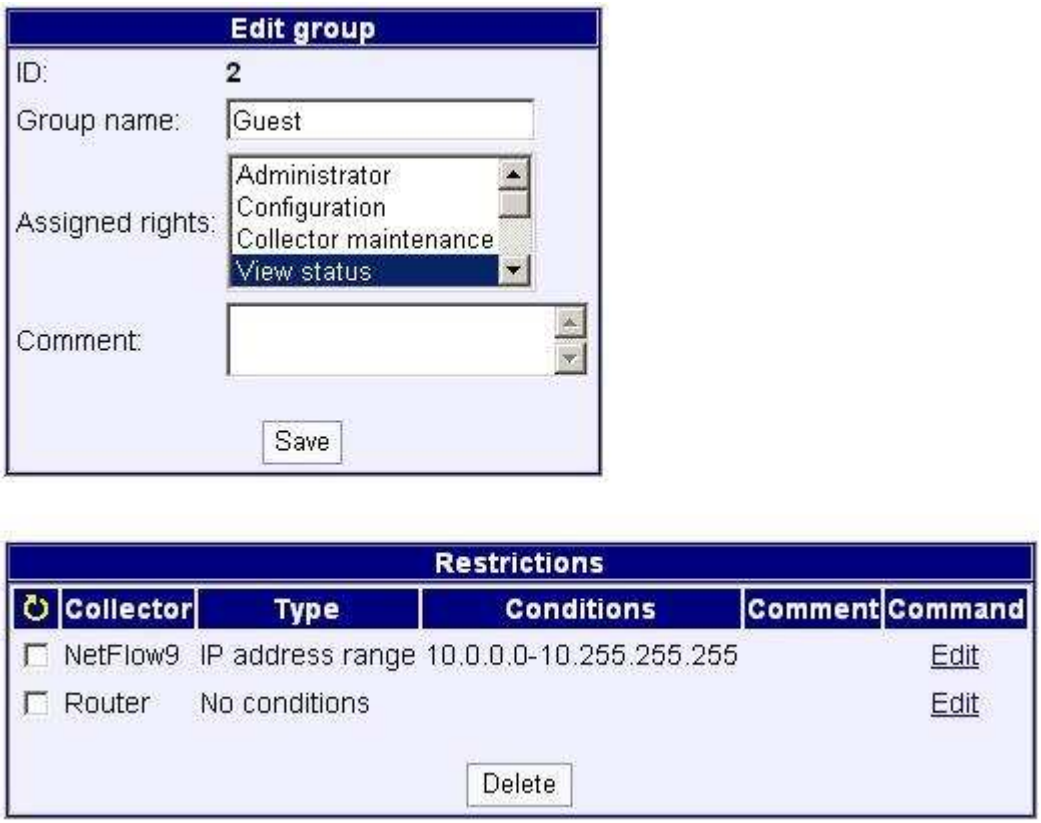

Figure: Group settings window.

User settings

In user settings you can create new users for the system. For each new user you will need to create a unique username. If the field password isn't empty, the user's password is changed to "typed new password". You can select a language, but the current version only supports English. In the next few months translations into French and German will be available. It is necessary to assign a user to the group. If the user account is disabled, select "enable account" by clicking on the item. Our software supports skins, so you can choose from several of our skins or define your own. All skins are saved in directory styles. The next option allows you to choose the size of the generated graphs and graph's colors. Allowed ranges for the graph x-axis is between 640 and 1800 and for the y-axis it is between 400 and 1600. The last option is for JavaScript support. We recommend using JavaScript extensions. The default system will automatically try to detect if JavaScript is enabled in your browser. If you have problems with JavaScript you can disable this extension.

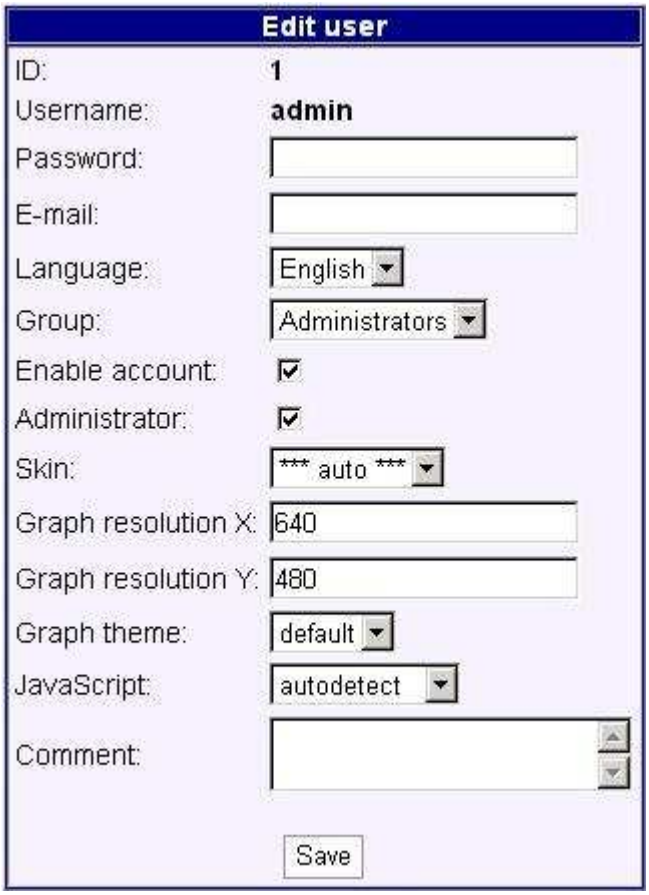

Figure: User settings window.

Account settings

Account settings are available for all users. In this menu you can change your password, select skins, graph sizes, etc. For more information about fields, read section "User settings".

If the global option "Display last logins" is enabled, you will see the ten last logins of your own below edit account window.

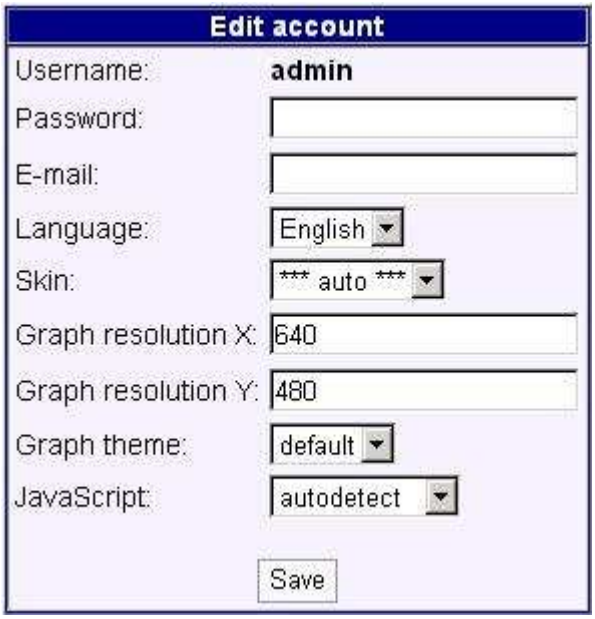

Figure: Account settings window.

5 User Guide

Main screen - Overview

After successful login, you will see the main screen dialog window. In the "Main screen" you will see all collectors, their state and some traffic statistics (Packets, Bytes and Flows counters). If you see any warnings, click on that link and find out what is wrong. Bad status is checked only for current hour.

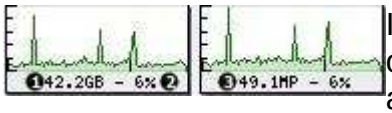

In the "Global setting" you can enable/disable displaying of utilization graphs on the overview page. The bytes (1) and packets (3) utilization histories are displayed under collector status. In each graph you can see its maximum value (1) and 5 minute utilization (2). The maximum utilization value is stored in the database for up to 1 month. Graph color depends on the utilization value (low value - green, middle value - yellow and high value - red).

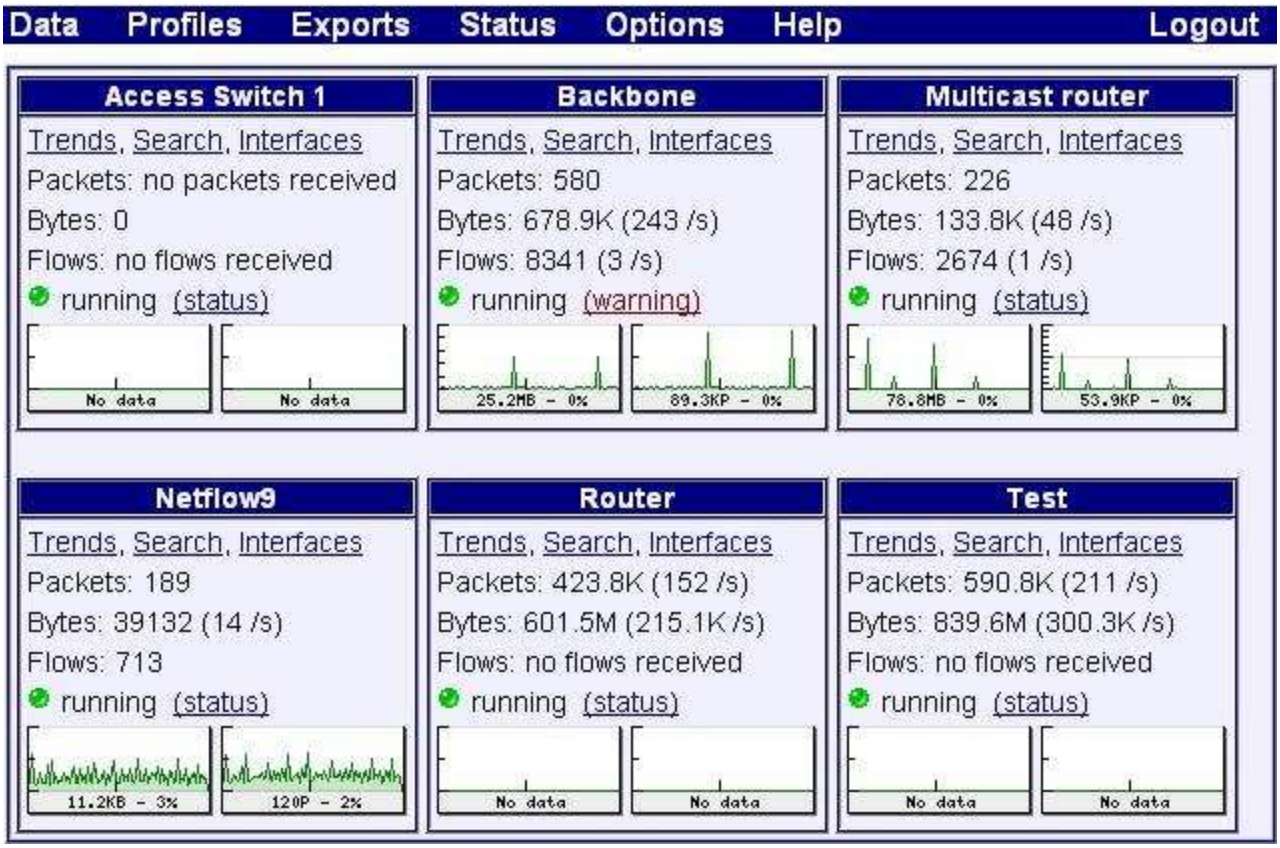

Figure: Main screen window.

You can select various items from the main menu:

- Data - traffic queries, information about IP addresses, graphs etc.
- Profiles - trends and search profiles.
- Exports - managing stored exports.
- Anomalies - view list of detected network anomalies.
- Status - state of engine, units, collectors and database.

- Options - configuration of this system.
- Help - documentation, license management, bug reporting etc.
- Logout - close session to web interface.

Data

In Data menu, there are main functions for traffic analysis.

- Overview - main screen window.
- Trends - many statistics, graph and table output.
- Search - detailed searching, output is formatted into table.
- Interfaces - input and output interface statistic, graph and table output.
- IP information - information about IP address (ping, whois etc).
- AS information - information about autonomous system from whois database.
- Graphs - displaying previously generated graphs via Trends menu.
- Utilization maps - managing and displaying of the utilization maps.

History

If you have enabled JavaScript functionality it's possible to use previously entered values in the dialog windows. If you would like to open a new history dialog window click on the icon located next to the selected field. The history dialog window will contain the last 30 entered values. The following window is an example of protocol history. If you want to clear the protocol history click on the "clear history" link.

Figure: History dialog window.

Trends

Trends are the most used menu in the whole system. This menu can run all wanted statistics. List of available statistics depends on selected table fields.

Trends conditions

To select table in "Table selector" first select the collector and then the table that you want to see. If you haven't enabled JavaScript, please, click on the "Select" button to choose the collector and then the wanted table. Your selection will be displayed in the information window below.

In "General parameters" first select one of the following statistic:

1. Bytes.
2. Packets.
3. Top source hosts per byte.

4. Top source hosts per packet.
5. Top source hosts distribution.
6. Top destination hosts per byte.
7. Top destination hosts per packet.
8. Top destination hosts distribution.
9. Top hosts conversations per byte. (*)
10. Top hosts conversations per packet. (*)
11. Top applications per byte.
12. Top applications per packet.
13. Top protocols per byte.
14. Top protocols per packet.
15. Top ToS/DSCP per byte.
16. Top ToS/DSCP per packet.
17. Top source TCP/UDP ports per byte.
18. Top source TCP/UDP ports per packet.
19. Top destination TCP/UDP ports per byte.
20. Top destination TCP/UDP ports per packet.
21. Top source interfaces per byte.
22. Top source interfaces per packet.
23. Top destination interfaces per byte.
24. Top destination interfaces per packet.
25. Top interface conversations per byte.
26. Top interface conversations per packet.
27. Top source ASes per byte.
28. Top source ASes per packet.
29. Top destination ASes per byte.
30. Top destination ASes per packet.
31. Top AS conversations per byte.
32. Top AS conversations per packet.
33. Top next hops per byte.
34. Top next hops per packet.
35. Top ICMP messages per byte.
36. Top ICMP messages per packet.

Note (*): If the statistic top conversations is chosen, domain name resolution is disabled in the graph.

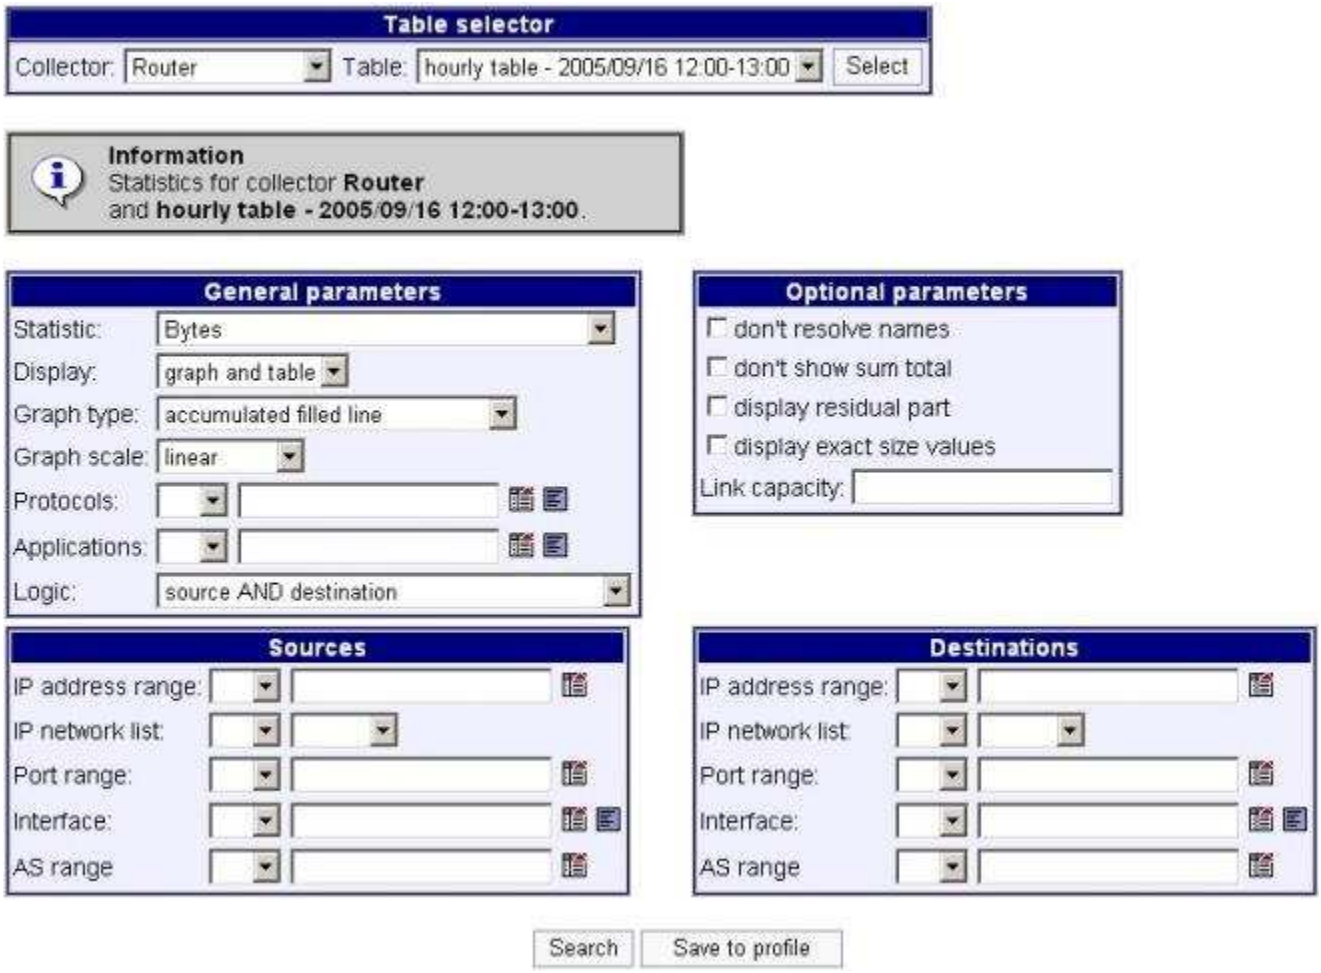

Figure: Specifying trends conditions.

The next options are related to formatting output, you can select if you want to generate a graph, table or both and what types of graph you want to see.

In the "time field" you can specify the time interval that you see. For example the tenth hourly table is: 10:20-10:45, and the weekly table is: 2006/02/15 – 2006/02/17. The list of times is separated by a comma. Click on the icon to display history window.

In the "bytes or packets field" you can specify which bytes or packets range you want to see. For example if you type in packet field value: 1 you will only see flows where only one packet is transferred.

In "protocols field" you can specify which protocols are seen. For example: TCP, UDP. The list of protocols is separated by a comma. A complete list of protocols is located in the system file /etc/protocols. Click on the icon to view list of defined protocols, applications or detected interfaces.

In applications field you can specify which applications you want to see.

Applications field can have the following formats:

- tcp/<portname> (e.g. tcp/smtp)
- tcp/<portnumber> (e.g. tcp/25)
- udp/same as for tcp (e.g udp/53, udp/domain)
- <protocolname> (e.g. gre, icmp, udp)
- <application_shortname> (e.g. dc). For application list see "Options->Applications".
- <application_number> (e.g. 300001). Check section Configuration, caption Applications is where ADM classify applications are described.

In "TCP flags" you can specify flags which you want to see. TCP flags field consists of one or two sets of characters <SAFRPU*> <SAFRPU*> separated by a space. Where character S stands for TCP flag synchronization, A for acknowledgment, F for finish, R for reset, P for push, U for urgent and * means all of the above. The first set of characters indicates which TCP flags must be set up, the second indicates which TCP flags you are checking.

Examples:

- SA * - find all flows with set up SYN and ACK flags, the remaining flags are not set
- SA SA - find all flows with set up SYN and ACK flags and ignore other flags
- S SF - find all flows with set up SYN flag and FIN flag is not set
- * - find all flows with set up all flags

If you enter only one set of characters (e.g. SA), the second is automatically set to "*".

The TOS byte in the IPv4 header has had various purposes over the years, and has been defined in different ways by five different RFCs (RFC 791, RFC 1122, RFC 1349, RFC 2474, and RFC 3168). The modern definition of the TOS byte is a six-bit Differentiated Services Code Point and a two-bit Explicit Congestion Notification field. For a full history of the TOS byte, see section 22 of RFC 3168.

Current CFI version accepts the following values:

• ToS values: 0-255
• DSCP values: AF11, AF12, AF13, AF21, AF22, AF23, AF31, AF32, AF33, AF41, AF42, AF43, BE, EF, CS1 - CS7, NC1 and NC2.
• RFC 791 specification: P0-P7DTR

where P0-7 means precedence value, character 'D' means minimize delay, character 'T' means maximize throughput and character 'R' means maximize relibility.

You can use arithmetic logic between source and destination window.

Possible values are:

1. source AND destination,
2. source OR destination,
3. source->destination OR destination->source.

In "Optional parameters" you can: disable domain names resolving, disable counting of total sums, enable displaying of residual part (residue of top ten), displaying exact size values (bytes instead of kilo or mega bytes equivalent) or convert byte values to the bits per second. You can specify link capacity that will be displayed in the graph. Link capacity is in the bits per second, but you can use values in kilobits or megabits, for example 10m means ten megabits per second.

Fields in source or destination windows can be different depending on the selected table.

The following are able to be viewed:

- IP address range (possible values):
  - Single IP address (10.1.1.1).
  - Domain name (web.mydomain.com).
  - List of IP addresses (10.1.1.1, 10.2.1.1, web.mydomain.com).
  - Range of IP addresses (10.3.1.1-10.3.255.255).
  - IP networks (10.0.0.0/8, 192.168.0.0/16).
  - IP network list defined via "Options->Networks".
  - Exclude range of network (10.0.0.0/8, !10.1.0.0-10.5.255.255)

  All previous types can be combined. Field separator can be comma or semicolon. You can also use an exclude character **!** which excludes single IP or range of IP from the list. Domain names can't be used when you use IP address ranges!

- IP network list. You can select network lists defined in "Options->Networks".
- Port range. In "Port" field you can use values that are same as those used in the "Applications" field but without application specific extensions (application short name or application number). (e.g. 80,135,137-139).
- Interface. You can use interface ifIndex number, list of interfaces or range (e.g. 1,10,20-25).
- AS range. You can use autonomous system number, list of autonomous systems or range (e.g. 1000,1902,5000-5005).

After completing the search conditions, you can start searching by clicking on the "Search" button or you can save search conditions in the trends profile by clicking on the "Save to profile" button. After saving conditions you will see information window (see picture bellow).

Figure: Saving conditions into profile.

Trends output

The pictures below show various examples of search results formatted into a graph.

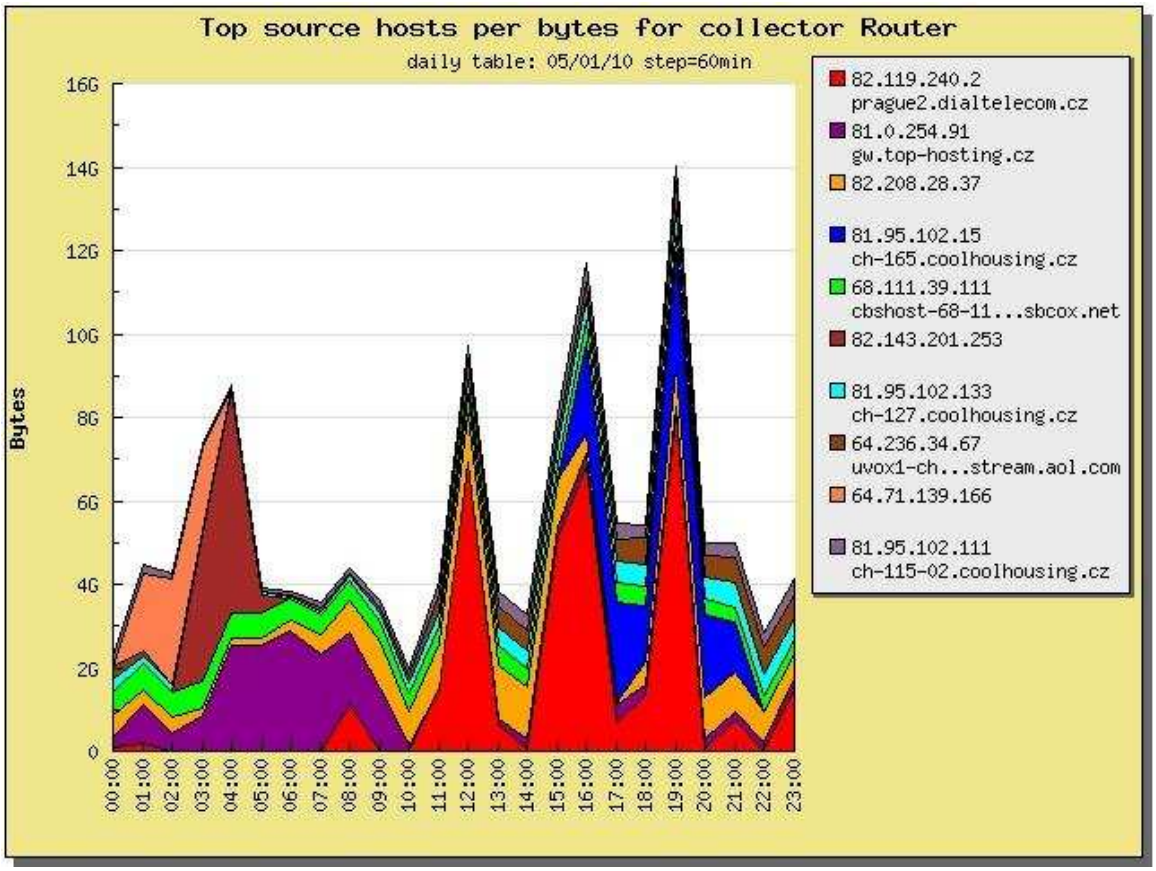

Figure: Accumulated lines graph.

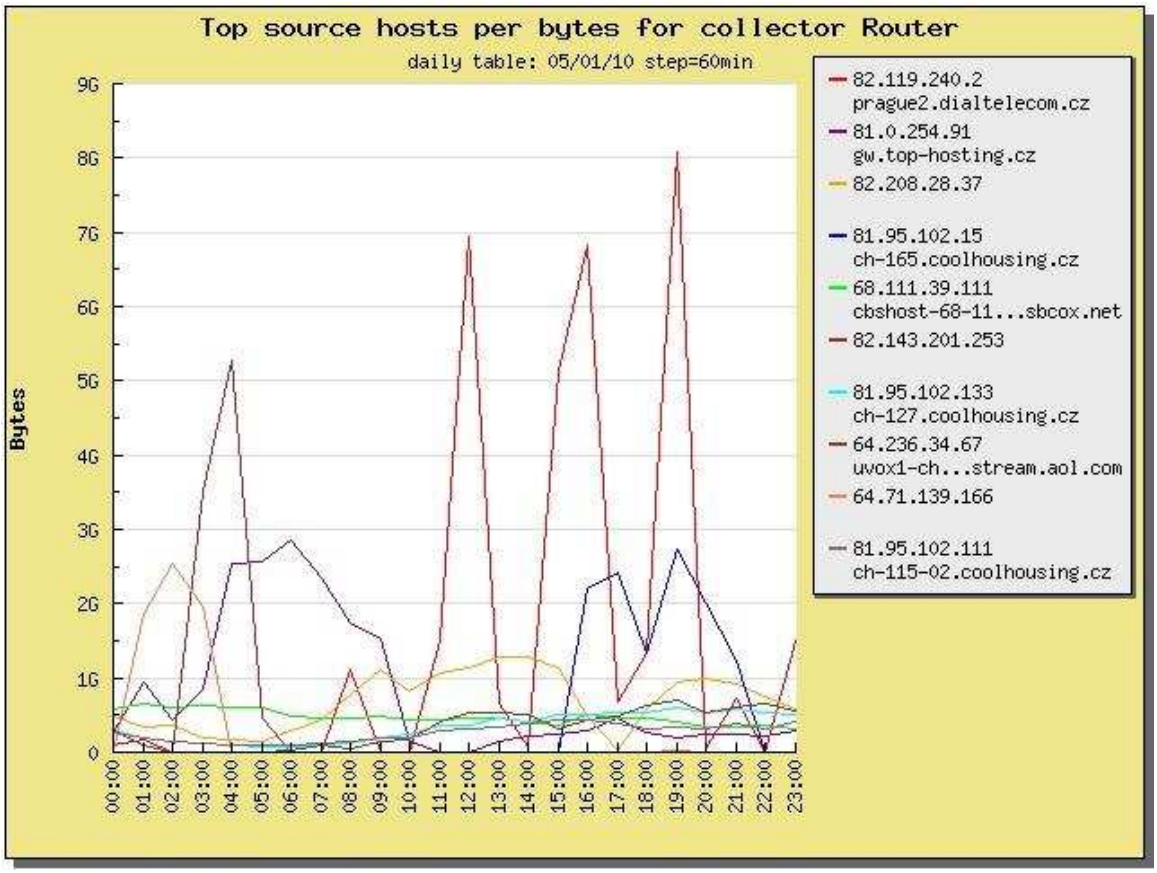

Figure: Non-accumulated lines graph.

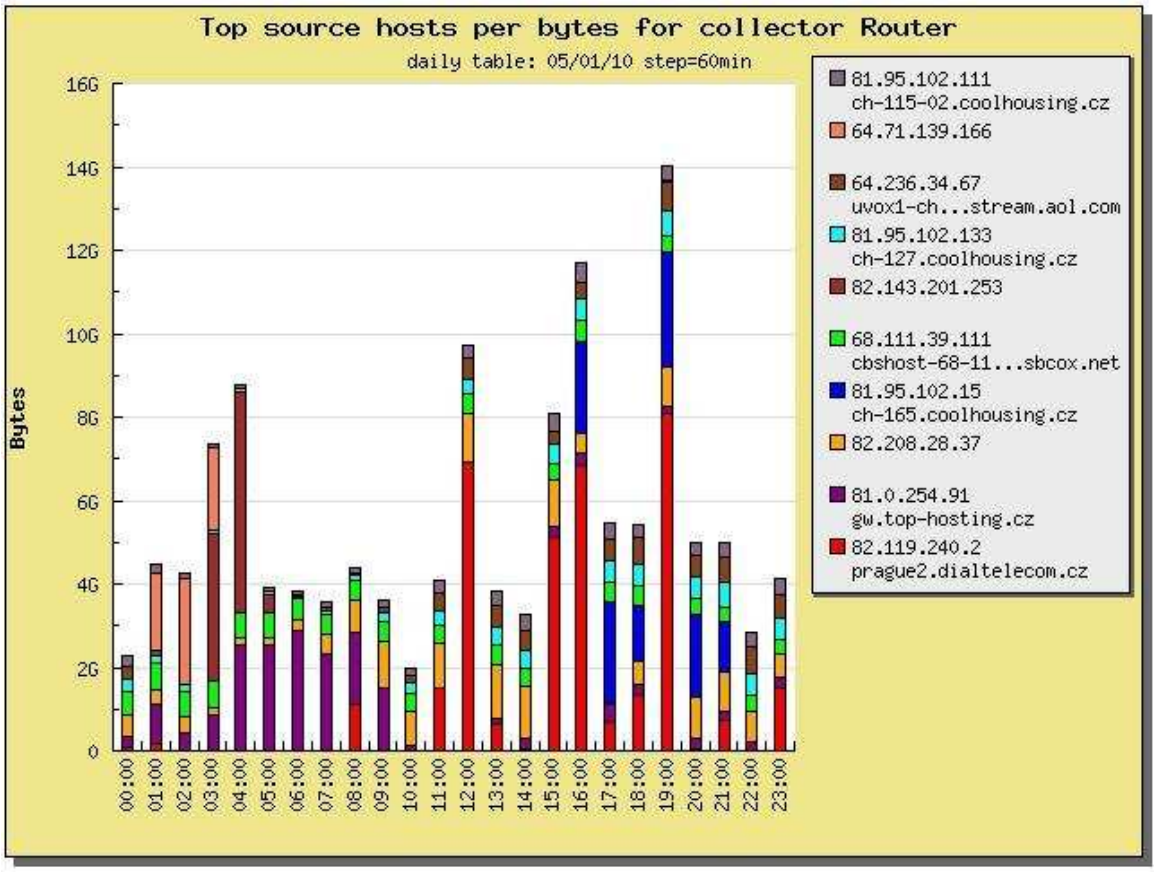

Figure: Accumulated bars graph.

This product offers various formats of search results. One of these options is format to table. An example of this is shown in the following picture:

Bytes statistics for collector	
Time	Bytes
Mar 21 2005 00:00:00	160.8G
Mar 21 2005 01:00:00	110.6G
Mar 21 2005 02:00:00	85.8G
Mar 21 2005 03:00:00	66G
Mar 21 2005 04:00:00	65G
Mar 21 2005 05:00:00	60.5G
Mar 21 2005 06:00:00	52.9G
Mar 21 2005 07:00:00	68.8G
Mar 21 2005 08:00:00	86.7G
Mar 21 2005 09:00:00	79.5G
Mar 21 2005 10:00:00	133.4G
Mar 21 2005 11:00:00	130G
Mar 21 2005 12:00:00	128.4G
Mar 21 2005 13:00:00	142.6G
Mar 21 2005 14:00:00	151.5G
Mar 21 2005 15:00:00	134.2G
Mar 21 2005 16:00:00	146.9G
Mar 21 2005 17:00:00	131.4G
Mar 21 2005 18:00:00	152.3G
Mar 21 2005 19:00:00	132.5G
Mar 21 2005 20:00:00	142.4G
Mar 21 2005 21:00:00	157.3G
Mar 21 2005 22:00:00	159.5G
Mar 21 2005 23:00:00	148G
Sum total	**2.8T**

Figure: Search results formatted into table.

Trends data export

Output data can be exported into CSV formatted file. This file can be opened in other applications for example in Microsoft Excel or in Open Office package. When you click on link "Export" in the left dialog menu, an export window will be displayed. You can then specify filename, time format and field header.

For time format you can use the codes listed bellow:

- %y - year as a decimal number without a century (range 00 to 99),
- %m - month as a decimal number (range 01 to 12),
- %d - day of the month as a decimal number (range 01 to 31),
- %H - hour as a decimal number using a 24-hour clock (range 00 to 23),
- %M - minute as a decimal number,
- %S - second as a decimal number,
- %Y - year as a decimal number including the century,
- %x - preferred date representation for the current locale without the time,
- %X - preferred time representation for the current locale without the date.

For example you can use time format: %x %X.

You can find a complete list of time formats in PHP documentation. Check web page: http://www.php.net/manual/en/function.strftime.php.

Export is saved into a temporary file. You can download this file via main menu "Exports". After successfully downloading it is recommended deleting this file to save disk space.

Trends email data

This feature allows you to send output data via SMTP protocol to a specific email address. When you click on the "Email results" link in the left dialog menu, an email window will be displayed. You can then specify an email address, subject and comment.

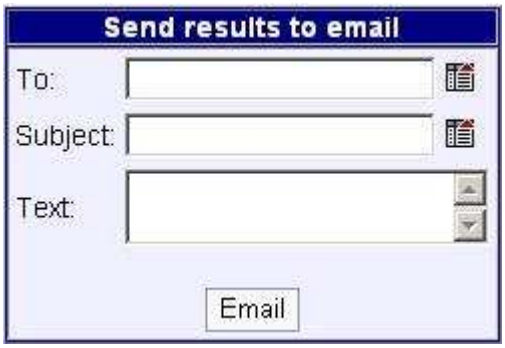

Figure: Email dialog window.

Search

In the "Search" menu (second most used menu) you can find detailed information about data flows. Output of search menu is always formatted into a table.

Search conditions

The "Search" menu contains a "Table selector" same as the "Trends" menu, its functionality is the same; see caption "Trends conditions" indicate how to manipulate the "Table selector". General parameters are nearly the same (without statistic list and graph format functions).

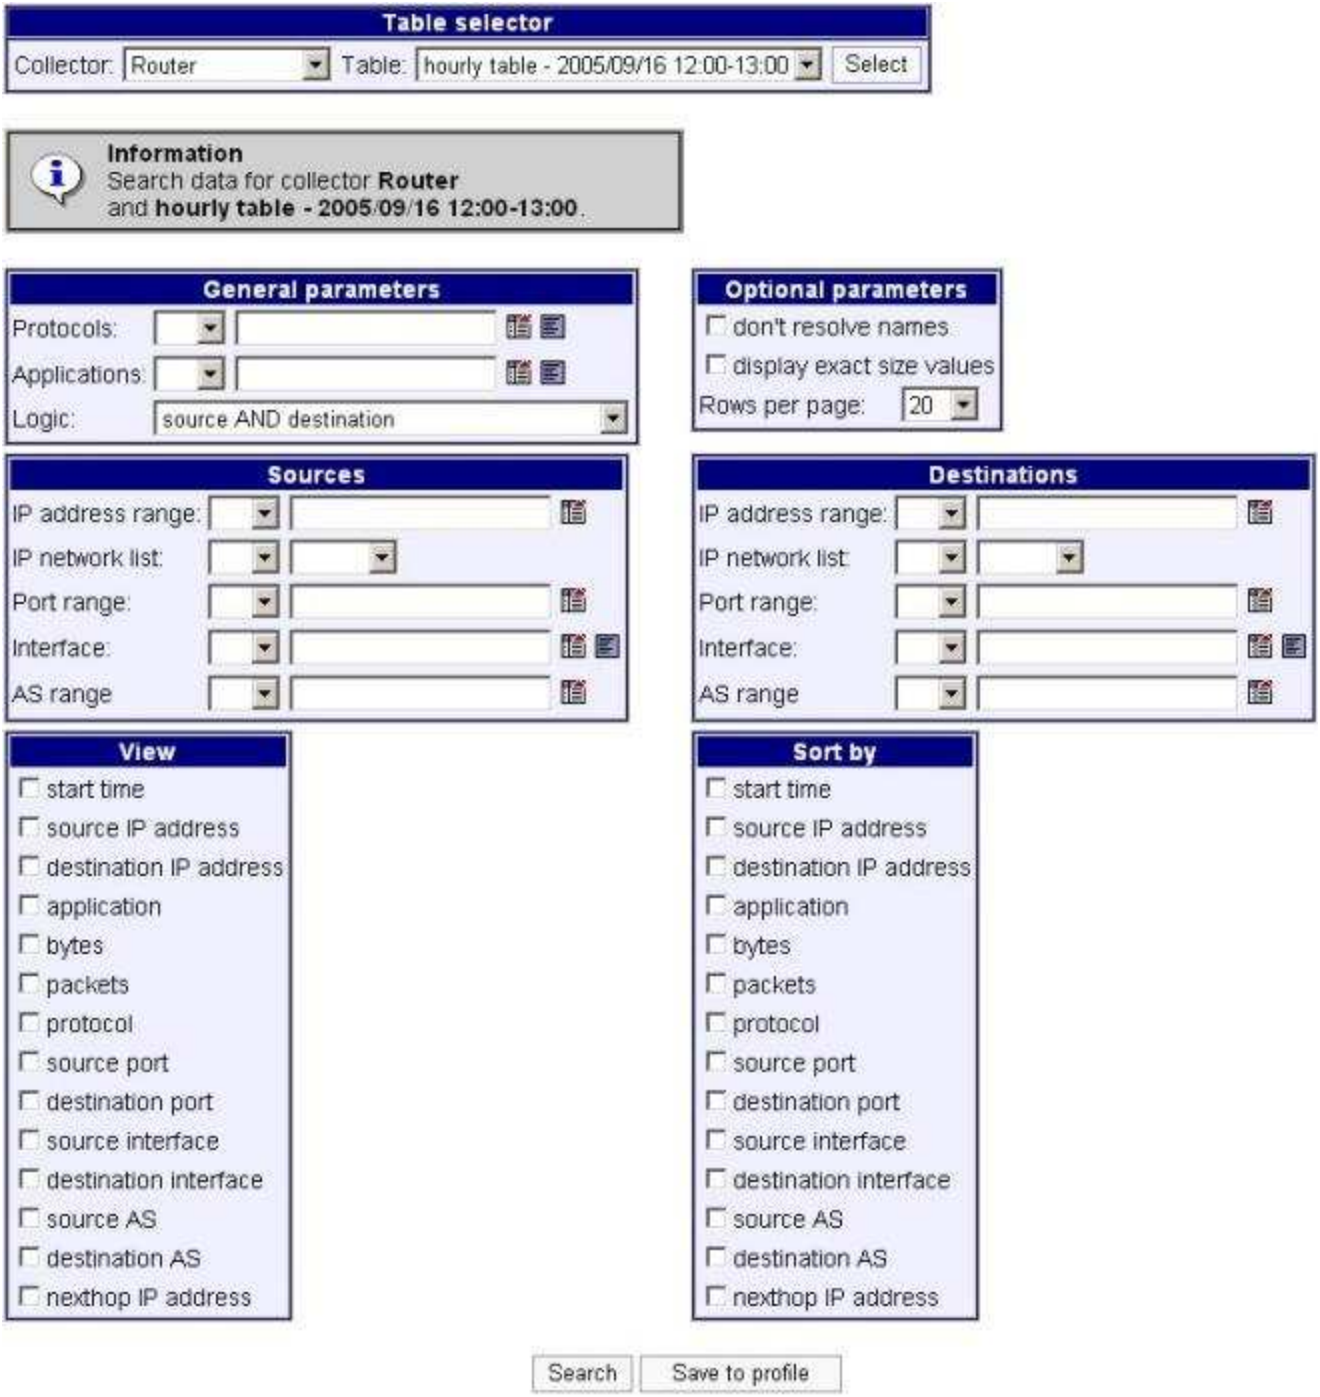

Figure: Search conditions.

In "Optional parameters" you can disable domain names resolution or change the number of lines per search result page (default 20 rows), you can change the number of rows per page to 100 maximum.

Fields in "Source" and "Destination" windows can change depending on the selected table. Please read the caption "Trends conditions" to get proper format of these fields.

The last two windows "View" and "Sort by" contain options for choosing which fields you want to see in the results. If you don't check any of the fields, by default, all fields are selected. For example you can select to see source and destination IP addresses, time and used application. If you want to see field bytes or packets, all other fields are grouped.

Search output

The picture below shows an example of search results formatted into a table.

Search result for collector Router hourly table: 05/03/22 10:00-11:00 step=1min					
Start time	**Destination IP address**	**Application**	**Bytes**	**Packets**	**Protocol**
Mar 22 2005 10:00:00	1.1.1.1	tcp/25 smtp	336	6	tcp
Mar 22 2005 10:00:00	4.14.52.27	tcp/42 nameserver	686	14	tcp
Mar 22 2005 10:00:00	4.15.9.110 wbar10...dsl-verizon.net	tcp/42 nameserver	898	6	tcp
Mar 22 2005 10:00:00	4.15.55.28 wbar2.a...dsl-verizon.net	tcp/42 nameserver	118	2	udp
Mar 22 2005 10:00:00	4.27.12.249	tcp/113 auth	144	3	tcp
Mar 22 2005 10:00:00	4.31.22.21 wbar100...dsl-verizon.net	tcp/42 nameserver	757	4	tcp
Mar 22 2005 10:00:00	4.60.62.218 lsanca1...dsl-verizon.net	tcp/25 smtp	322	7	tcp
Mar 22 2005 10:00:00	4.62.161.53 lsanca2...dsl-verizon.net	tcp/25 smtp	1.3K	7	tcp
Mar 22 2005 10:00:00	4.68.245.3 msnntm1b.level3.net	tcp/42 nameserver	1K	2	udp
Mar 22 2005 10:00:00	4.68.245.66	tcp/25 smtp	137.9K	1.7K	udp
Mar 22 2005 10:00:00	4.78.22.9 ns1.techiemedia.net	udp/53 domain	270	1	udp
Mar 22 2005 10:00:00	4.78.57.166	tcp/80 www	92	2	tcp
Mar 22 2005 10:00:00	4.234.21.189 dialu...miami1.level3.net	tcp/42 nameserver	772	12	tcp
Mar 22 2005 10:00:00	4.253.98.152 dialup-4.253...level3.net	tcp/42 nameserver	92	2	udp
Mar 22 2005 10:00:00	8.7.146.108 zztesa.com	tcp/25 smtp	60	1	tcp
Mar 22 2005 10:00:00	9.254.55.28	tcp/80 www	96	2	tcp
Mar 22 2005 10:00:00	10.0.0.66	tcp/22 ssh	96	2	tcp
Mar 22 2005 10:00:00	10.0.1.128	tcp/153	96	2	tcp
Mar 22 2005 10:00:00	12.4.198.133 cws3.crutchfield.com	tcp/80 www	16.9K	291	tcp
Mar 22 2005 10:00:00	12.15.136.40	tcp/80 www	697	5	tcp

results 1 - 20

Figure: Search query result.

If the result of a search query is more lines than maximum rows per page value, you can click on "Next" button to see the next page. If you want to redefine query, click on the "REDEFINE" link in the left menu. Click on the "NEW QUERY" link for blank search condition dialog. In the "Search" menu there are functions to save search conditions to profile (see "Trends" caption), export output to CSV file (see "Trends

export data" caption) or send output to email address (see "Trends email data" caption).

Interfaces

The "Interface" menu (third most used menu) contains information about device interface utilization.

Interface conditions

The "Interface" menu contains a "Table selector" which is the same as the one found in the "Trends" menu. Its functionality also the same; see caption "Trends conditions" this shows how to manipulate the "Table selector". You can use this menu only for the tables with the source or destination interface index fields. General parameters are nearly the same (without statistic list). In the interface item you can also specify which interface you want to apply the statistic (E.g. 1,5-8). If selected collector is associated with more SNMP enabled devices, you can specify an interface in the format:
#device1_ip_address:interface_index,#device2_ip_address:interface_index
(E.g. #10.1.1.1:1-8,#10.1.2.1:5-20).

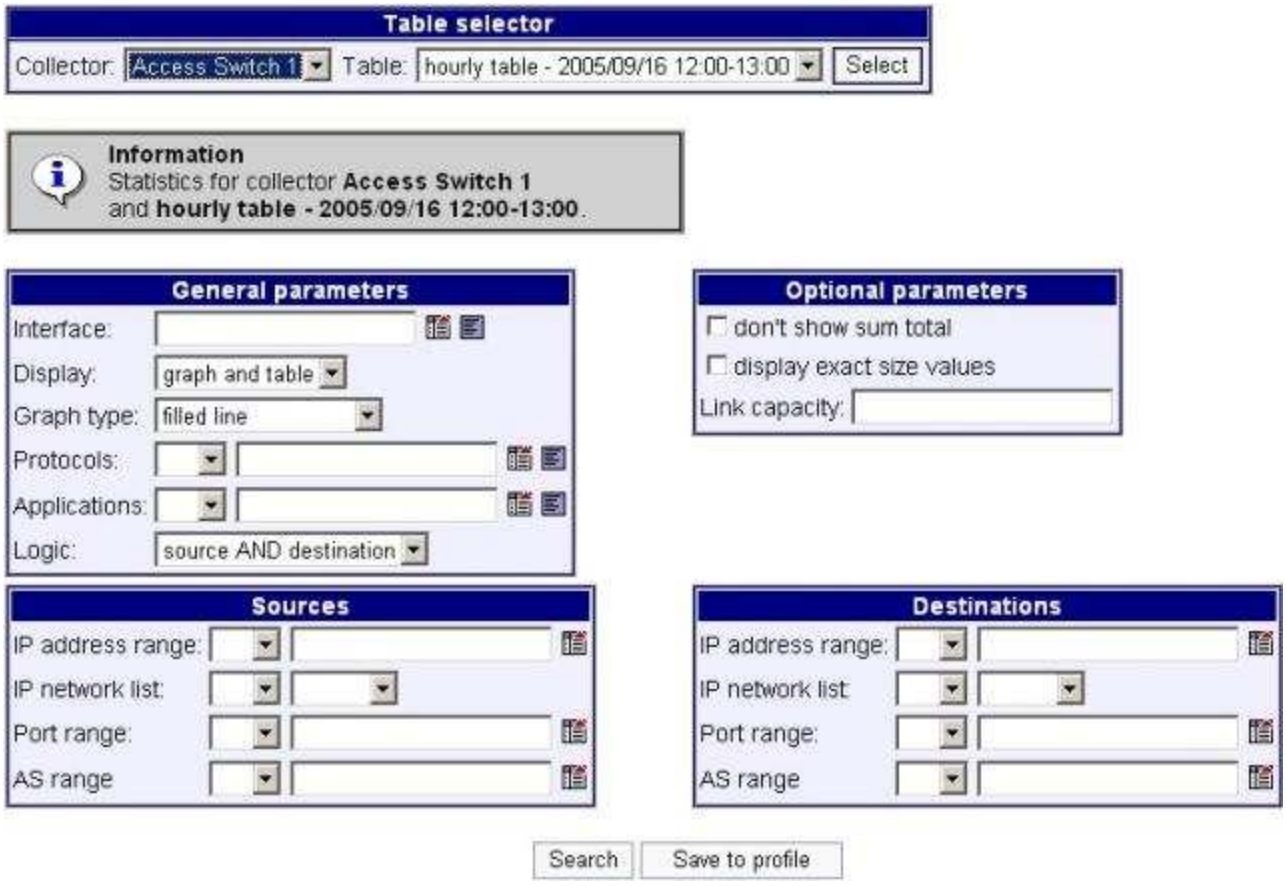

Figure: Interface conditions.

In "Optional parameters" you can disable the counting of total sums or displaying exact size values (bytes instead of kilo or mega bytes). You can specify link capacity

that will be displayed in the graph. Link capacity is in bits per second, but you can also use values in kilobits or megabits (i.e.10m means ten megabits per second).

Fields in "Source" and "Destination" windows can change depending on the selected table. Please read the caption "Trends conditions" to get proper format of these fields.

Interface search output

The picture below shows an example of interface statistics.

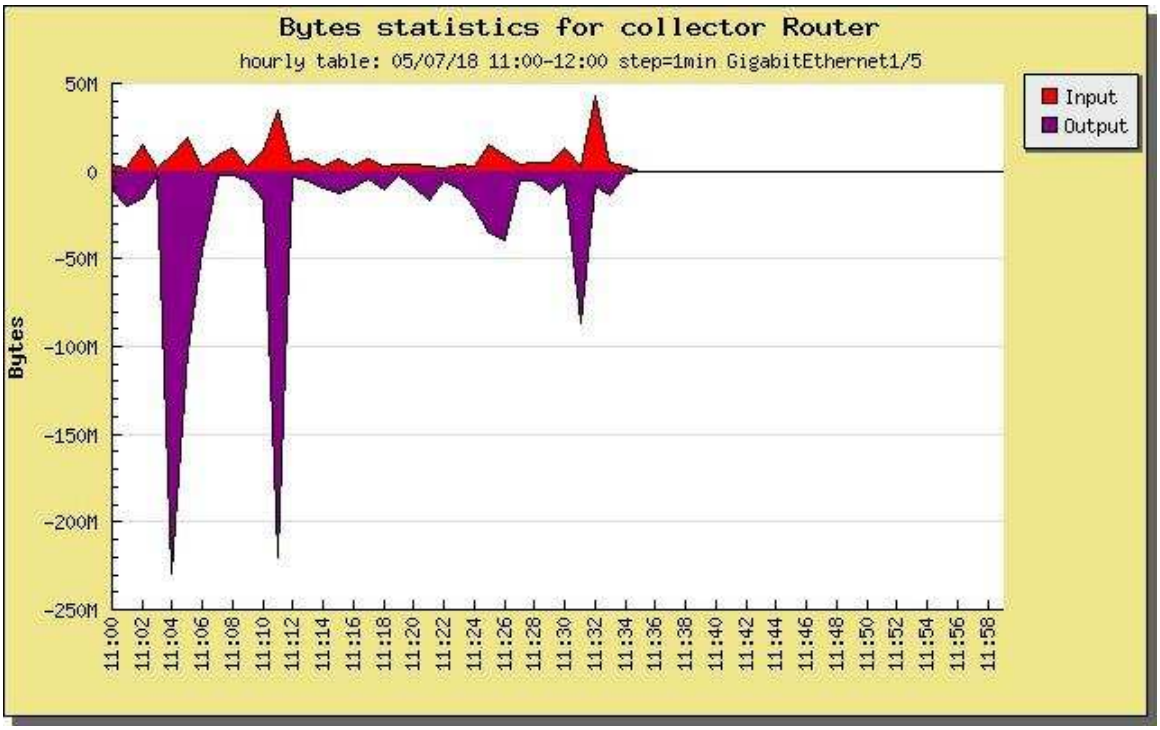

Figure: Interface statistics results.

If you want to redefine the query, click on the "REDEFINE" link in the left menu. Click on the "NEW QUERY" link for a blank search condition dialog. The "Interface" menu contains functions that save interface conditions to profile (see "Trends" caption) and exports the output to a CSV file (see "Trends export data" caption). You can also send the results via email (see "Trends email data" caption).

IP information

The "IP information" menu contains functions for getting information about used IP address(es). This option gives you the possibility to see domain names (if one exists), IP address class (in a classful network), country and autonomous system related information. IP address to country or autonomous system mapping can be changed in the "Options->Country" menu or in the "Options->AS list" menu. If you have rights to run shell commands, you can ping the IP address and trace the route to it's destination, query whois database or try querying HTTP server using the HTTP HEAD method.

Figure: Basic IP address information.

AS information

In "AS information" menu you can query whois database to get information about autonomous system. In default whois server is determined automatically, but you still have the ability to specify which server you want use.

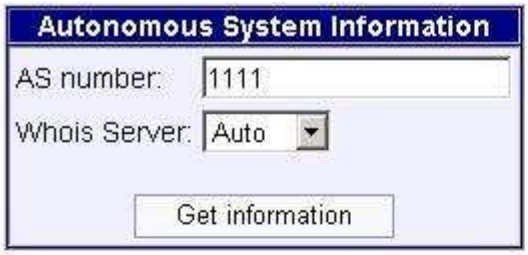


```
% This is the RIPE Whois query server #2.
% The objects are in RPSL format.
%
% Rights restricted by copyright.
% See http://www.ripe.net/db/copyright.html

as-block:       AS1101 - AS1200
descr:          RIPE NCC ASN block
remarks:        These AS numbers are further assigned by RIPE NCC
remarks:        to LIRs and end-users in the RIPE NCC region
remarks:        Please refer to these documents
remarks:        <http://www.ripe.net/ripe/docs/ir-policies-procedures.html>
                <http://www.ripe.net/ripe/docs/asnrequestform.html>
                <http://www.ripe.net/ripe/docs/asnsupport.html>
org:            ORG-NCC1-RIPE
admin-c:        CREW-RIPE
tech-c:         OPS4-RIPE
mnt-by:         RIPE-NCC-HM-MNT
mnt-lower:      RIPE-NCC-HM-MNT
changed:        er-transfer@ripe.net 20020822
changed:        ripe-dbm@ripe.net 20040421
source:         RIPE
```

Figure: Basic autonomous system information.

Graphs

All graphs generated by "Trends" menu are saved for later viewing. System saves these images for one day (images have a cache flag). In "Graphs" menu you can view these images or save cached images (set flag to saved value). With save flag, graph will not be deleted after one-day timeout. You can view all images or just the selected one. User with administrator right can see the images of any other user.

Figure: List of stored graphs.

Utilization maps

In the menu “Utilization maps” you can define maps with one or more objects and paths. For every object you can define certain conditions (e.g. IP address networks). Caligare Flow Inspector will count 5 minute byte utilization for each object and display the results on the public available map. This map (or simply image) can be linked from any other web page. For example you can define maps for displaying utilization of web services, FTP transfers or overall network activity. Click on the “View” link to display utilization map with measured values. Generated pictures are cached; click on the “Clear cache” link for clearing cached image.

Figure: Example of utilization map.

Creating new utilization map

Add new map if you want to use utilization maps. First select the collector that you want to count the utilization statistic (collector can’t be changed later). Next the parameters are map name and map caption. Map name is required and must be unique. Map caption is optional, if this field is filled up; caption is displayed on top of the image. If you want to enable generating utilization map, activate map by checking the “Active” box. The next required option is background image. Select one of the images in the list or simply click “Viewer” button, which is a background image wizard tool.

If you enable the option “show collector name” the collector’s name is displayed in the bottom left portion of the image. The option “show measurement time” enables displaying 5 minute time interval measurements in the bottom right portion of the image.

The final two options are for selecting colors. The first is “Text color” – which selects the color of the texts in the image, the second is the “Path color” – which chooses color of the paths connecting the objects.

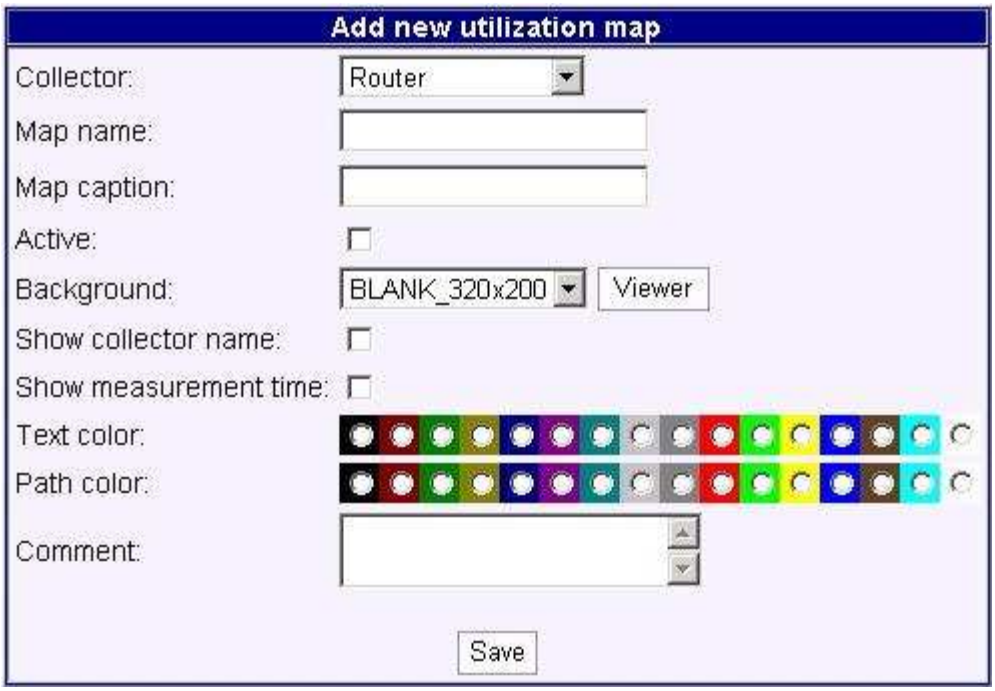

Figure: A new utilization map dialog window.

Utilization objects

To view objects associated with map click on “Objects” label in the utilization maps list. If you want to change the position of the object; click on “Position” label in the utilization objects list. Set up a new object position by clicking on the map area, where you want to place your new object. Position is stored in database automatically.

New object can be created only if collector is running and has valid hourly tables. For each object you can specify many options. Standard parameters are parameters such as object name and object caption. Object name is required and must be unique for the selected map. Object caption text is optional and if this field is fill up; caption is displayed on top of the object image. Next two parameters are for utilization. Parameter capacity is 5 minute of traffic that is transferred through object. Capacity value is number of bytes per 5 minutes interval. For example you can enter a value of 10M, the object is transferred 10 megabytes/5 minutes in the peak hours.

If you don’t know the exact value, check only the second parameter “Capacity auto-learn”. This parameter will store object peak utilization every time you view the utilization map. For each object you can select an image. Select one of the images in the list or simply click on “Viewer” button, which is an object image wizard tool. The

following two parameters are related to image position. If you don't know the position pixels, after saving the object, click on "Position" label in the list of objects to set it up. Remaining options can be different depending on collector settings. Check chapter "Trends conditions" for detail.

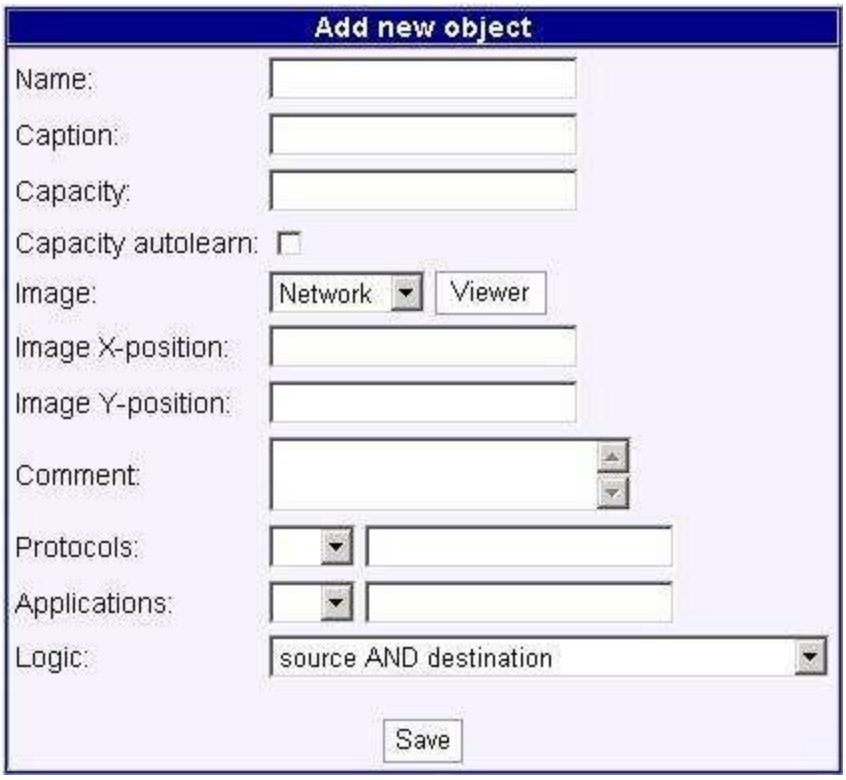

Figure: A new utilization map - objects dialog window.

Utilization paths

To view a list of paths associated with the map click on "Paths" label in the utilization maps list.

New path can be created only if you have a minimum of two objects. Selection of unique source and destination object is required. You can set up the weight of the path between 1 and 6. For example you can use weight 1 for dial-up lines, 2 for serial lines, 3 for 10Mb/s lines, 4 for fast Ethernet lines, 5 for Giga Ethernet lines and 6 for 10G Ethernet lines.

Figure: A new utilization map - paths dialog window.

Profiles

The "Profiles" menu allows you to manage stored trends and search profiles.

Figure: List of stored global trends profiles.

Each profile has a global or local flag. Profiles with a global flag are available for all users; those with a local flag are available only for user who saves it. Click on "Edit" link if you want to modify profile name or flag. For changing trends or search conditions, click on label "Modify conditions" in window "Edit profile". "Trends" or "Search" menu will be displayed; edit conditions that you want to change and click on the "Save to profile" button. The selected profile will be replaced.

Figure: Edit trends profile.

Search profiles have the same functionality as trends profiles.

Figure: List of stored user's search profiles.

Exports

Export list

When you export rows from "Trends" or "Search" menu they will be saved into a temporary file. This file can be downloaded via "Export list" menu. Click on filename to download the selected file to your computer. We recommend deleting exports after downloading to save your disk space. Free space on your server is displayed in the informational window above list of available exports. Users with administrator right

can see exports of all the users. In the export list you can find exported whole data tables. Exported data tables can be imported later.

Figure: List of exported files.

Export status

If you request export data table (in the "Status->Tables" menu) the request will be queued and import/export daemon will dump these tables in 15 minutes. This menu allows you to see request queue and state of export process. If the import/export process doesn't start in 2 hours after inserting request, a warning window will be displayed.

Import list

In the "Import list" menu are exported files that can be imported into the system. Imported tables are standalone tables and the collector process can't remove them. If you want to import tables simply select the table and click on the "Import" button. State of import process shows the "Import status" menu.

Import status

Requesting the import data table will be queue the request and the import/export daemon will insert these tables in 15 minutes. If the import/export daemon doesn't start in 2 hours from the inserting request, a warning window will be displayed. Imported table will be available for statistics only when the collector process is running and finish aggregating any other tables.

Anomalies

You can view list of network incidents in the 'Anomalies menu'. Every incident consists of one or more alerts; every alert consists of one or more anomalies. You can use list filtering, i.e. by time range (current hour, last 3 hours, 24 hours, 3 days, 7 days, 14 days, 31 days or any time), by severity (only critical, urgent or higher, important or higher, warning or higher, or any severity), by related network (only internal networks, only external networks, any network), and by state (new, solving, resolved, archived or any state).

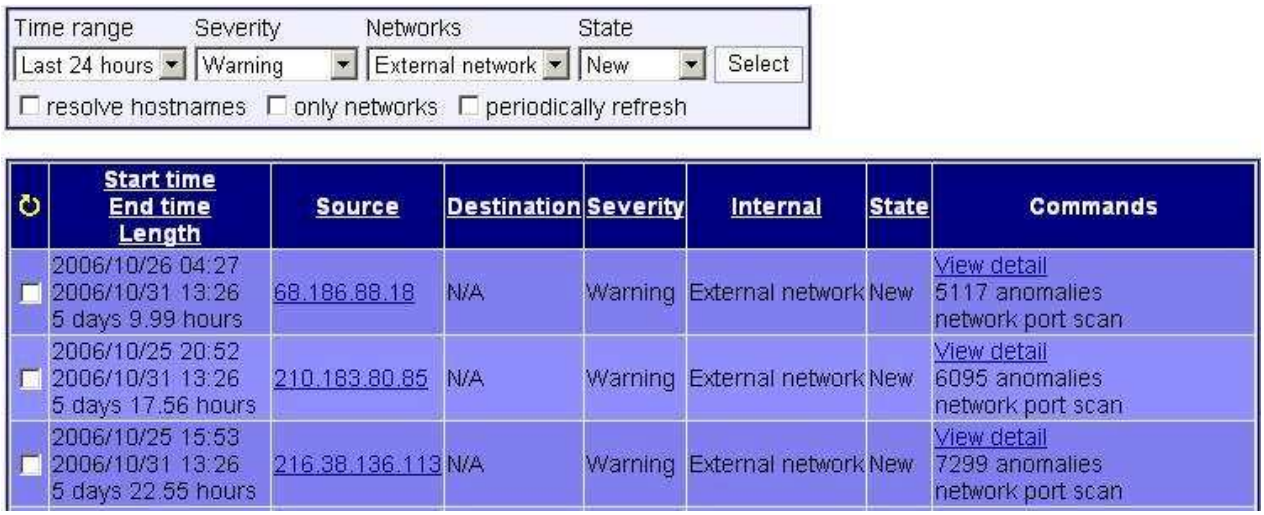

Figure: Anomalies window.

Host name resolving is disabled in default setting, but you can click on 'resolve hostnames' to receive full hostname. The other option gives you the possibility of viewing only network groups instead of full hosts IP addresses. The last option is used for periodical refreshing of the selected page. If you enable this option page it will be refreshed in 1 minute interval. This option is very often used by network security operators.

A list of detected network incidents is available below the filter window. You may order rows by clicking on the field header (click for the second time to descendant order).

By clicking on the source and/or destination (if available) you will receive a list of available actions for each row (i.e. anomaly exclusion, IP address information, more deep searching via Data->Search etc.). To view incident detail (and a list of alerts) click on the 'View detail' link. In the 'List of alerts' you may type in your comments, set state of incident, report incident to the email address or view anomaly details. See chapter Configuration – Anomalies and Appendix 3 for more information about network anomalies.

Status

In the "Status" menu you can get information about state of all system components. To view information about any component, click on one of following links:

- Engine - state of installed components.
- Devices - state of devices, list of interfaces, link error numbers.
- Units - state of units, display units running processes.
- Collectors - state of collectors, number of received flows etc.
- Last login - list of all software login.
- Tables - list of all flow tables, sizes, number of rows etc.
- Database - list of running database processes, list of all database tables.

Engine

"Engine" submenu shows the state of PHP, SNMP, graphic library and database library. If all components are functional in "Status" they will be installed.

Component	Status	Description
PHP:	installed	4.3.10-2
SNMP:	installed	
Graphics library:	installed	GDlib version 2
Database library:	installed	server: 4.0.23_Debian-3-log client: 3.23.56

Figure: State of installed components.

Devices

In "Devices" menu there is a list of configured devices (see picture bellow).

Name	Comment	Command
Access Switch 01		Detail
Backbone Router 01		Detail

Figure: List of devices.

If an IP address and/or SNMP community is configured you can see detailed device information.

System information	
System name:	
Description:	
Location:	
Contact:	
Uptime:	
Services:	

ifIndex	Description	Type	MTU	Speed	Physical address	Admin status	Operation status	Input octets
1	GigabitEthernet0/1	ethernetCsmacd	1500	1G		up	up	817M
2	GigabitEthernet0/2	ethernetCsmacd	1500	10M		down	down	0
3	GigabitEthernet0/3	ethernetCsmacd	1500	1G		up	up	3G
4	GigabitEthernet0/4	ethernetCsmacd	1500	1G		up	up	1G
5	GigabitEthernet0/5	ethernetCsmacd	1500	1G		up	up	1G
6	GigabitEthernet0/6	ethernetCsmacd	1500	1G		up	up	3G
7	GigabitEthernet0/7	ethernetCsmacd	1500	1G		up	up	373M

Figure: Detail device information.

Units

In "Units" submenu you can check the state of all configured units. Before each unit name is a displayed LED indicator. Green indicator means that the unit process is running and the unit is ready to manage collectors. If the red indicator is displayed the unit can't run and will not communicate with the database or communication between unit server and database server is unsynchronized. In order to synchronize

server's time we recommend using an *ntpdate* package. To resolve other problems see the "Installation section".

State	Name	Command
	localhost	Detail
	second machine	Detail

Figure: List of configured units.

Click on "Detail" link to get more detailed information about processes that use unit master process.

Unit - localhost	
Type	**Data on 2005-09-16 12:17:50**
System load:	72% 88% 98%
Process ID (pid):	8910
Start time:	2005-09-15 13:31:22
Uptime:	22 hours, 46 minutes and 37 seconds

Sub Processes			
Type	**ID/port**	**Process ID (pid)**	**Last echo at**
Aggregation	0	8912	2005-09-16 12:17:59
Database	0	8911	2005-09-16 12:17:48
Worker	33333	8914	2005-09-16 12:17:56
Worker	7777	8916	2005-09-16 12:17:51
Worker	33334	8913	2005-09-16 12:17:50
Worker	60000	8915	2005-09-16 12:17:54

View collectors used by current unit

Figure: Detail unit information.

Collectors

In "Collector status" you can see the state of all configured collectors. In front of the collector name is an LED indicator. Green LED indicator means that the collector is running, red LED means that the collector is disabled and a blinking red LED means that the collector is enabled, but not running. If a unit is ready, but the collector still doesn't run, see syslog messages on the unit server for error messages. A non-running collector is indicated after 30 seconds of inactivity.

State	Name	Command
	Access Switch 1	Detail
	Coll1	Detail
	Multicast router	Detail
	Router	Detail

Figure: List of configured collectors with their states.

Click on "Detail" link to get more information about a specific collector. "Detail" link gives you detailed information about collector process start time, current hour and summary statistics (number of received packets, bytes and flows, forwarded and dropped packets etc). Zero number of received packets may signify data link problems or a badly configured export device. DoS value and state, are values that indicate the state of denial of service protection. If a non-zero value is stated, DoS protection is activated. See syslog on the unit server for more information.

Collector Multicast router	
Start time:	2006-07-14 11:42:43
Uptime:	5 days, 23 hours, 49 minutes and 32 seconds

Collector Multicast router (current hour)		
	Suma	Avg. per sec
Number of packets:	153	0.1
Number of bytes:	86616	44.8
Number of flows:	1728	0.9
Forwarded packets:	153	0.1
Dropped packets due to bad source IP:	0	0
Dropped packets due to unsupported netflow version:	0	0
Dropped flows due to corrupted time:	0	0

Collector Multicast router (summary)		
	Suma	Avg. per sec
Number of packets:	41751	0.1
Number of bytes:	24562776 (24.6M)	47.4
Number of flows:	490849 (490.8K)	0.9
Forwarded packets:	41750	0.1
DoS protection state:	0	0
Dropped packets due to bad source IP:	0	0
Dropped packets due to unsupported netflow version:	0	0
Dropped flows due to corrupted time:	0	0

Figure: Detailed collector information.

Last login

In the "Last login" menu only the user with administrator rights can see who loged in to the web interface. Only last 300 logins are displayed. If you enable global option "Display last logins" each user can see last ten logins in the menu "Options->Account".

Date	Username	From
2005/09/16 12:11	admin	127.0.0.1
2005/09/16 10:13	admin	127.0.0.1
2005/09/16 09:53	admin	127.0.0.1
2005/09/16 09:34	admin	127.0.0.1
2005/09/15 16:14	admin	127.0.0.1
2005/09/15 14:46	admin	127.0.0.1
2005/09/15 14:03	admin	127.0.0.1
2005/09/15 13:33	admin	127.0.0.1
2005/09/15 09:39	admin	127.0.0.1
2005/09/14 17:33	admin	127.0.0.1
2005/09/14 16:15	admin	127.0.0.1
2005/09/14 13:42	admin	127.0.0.1

Figure: Last login information.

Tables

"Tables" menu transparently shows a list of used flow tables. This list of used tables may be very large. To view used tables for the selected collector, select a table(s) by click the selection box. If JavaScript is disabled, click on the "Select" button. If that table has a flag (previous, actual, next or moving data table) it cannot be deleted at this moment. If you want to export tables, select them and clink on the "Export" button.

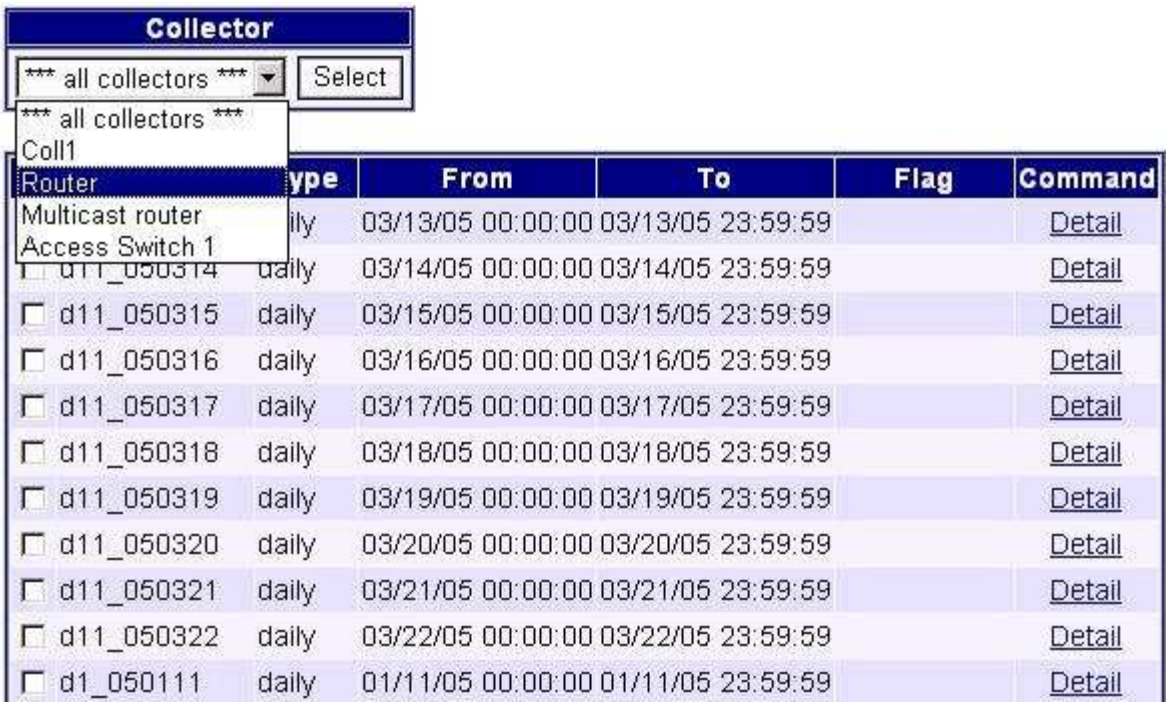

Figure: List of used flow tables.

Click on the "Detail" link to see how many rows are in the table, data and index sizes and when the table is aggregated. You can also see which tables are aggregated into the selected table.

Name:	**d2_050321**
Type:	daily
From:	03/21/05 00:00:00 (1111359600)
To:	03/21/05 23:59:59 (1111445999)
Flag:	
Rows:	1397516
Data length:	64285736 (64.3M)
Index length:	67485696 (67.5M)
Total length:	131771432 (131.8M)
Including:	h2_05032100 h2_05032101 h2_05032102 h2_05032103 h2_05032104 h2_05032105 h2_05032106 h2_05032107 h2_05032108 h2_05032109 h2_05032110 h2_05032111 h2_05032112 h2_05032113 h2_05032114 h2_05032115 h2_05032116 h2_05032117 h2_05032118 h2_05032119 h2_05032120 h2_05032121 h2_05032122 h2_05032123

Figure: Detail flow table information.

Database

"Database" menu is used to check database status. If the database is very loaded, some of counters may overflow. All running database threads can be viewed from the "Processes" submenu. If you want to stop a long-running query, when in your browser, (query continues to run and consume processor time), database menu allows you to kill long running threads. Starting a long-run query ("Search" or "Trends" menu) may cause problems, because the web server or PHP can close the connection. You can view all of the data tables by clicking on the "Tables" link.

User	Host	DB	State	Time	Info	Command
root	localhost	nfx	Sleep	1		Kill thread
root	localhost	nfx	Sleep	7		Kill thread
root	localhost	nfx	Sleep	2		Kill thread
root	localhost	nfx	Sleep	3		Kill thread
root	localhost	nfx	Sleep	4		Kill thread
root	localhost	nfx	Sleep	8		Kill thread
root	localhost	nfx	Sleep	2		Kill thread
root	localhost	nfx	Sleep	38		Kill thread
root	localhost	nfx	Sleep	11		Kill thread
root	localhost	nfx	Sleep	38		Kill thread
root	localhost	nfx	Sleep	37		Kill thread
root	localhost	nfx	Sleep	38		Kill thread
root	localhost	nfx	Sleep	37		Kill thread
root	localhost	nfx	Sleep	37		Kill thread
root	localhost	nfx	Query	0	SHOW PROCESSLIST	Kill thread
root	localhost	nfx	Sleep	37		Kill thread
root	localhost	nfx	Sleep	37		Kill thread
DELAYED		nfx	Delayed_insert	Waiting for INSERT	19	Kill thread
DELAYED		nfx	Delayed_insert	Waiting for INSERT	4	Kill thread

Figure: List of running database processes.

Options

"Options" menu is described in caption Configuration.

Help

In "Help" menu you can find functions for getting information product version, about PHP configuration, TCP and UDP ports and managing license.

Port database

In "Port database" there is a list of some well-known ports. You can get detailed information about a used port by clicking on the port number. This database is being continuously updated. The informational window shows you more detailed information about known problems, descriptions, server and client programs that are using this protocol and URL address.

results 1 - 100

Port number	Name	Description
0	none	
1	tcpmux	TCP multiplexer
2	compressnet	defunct service and a trojan
2	death	defunct service and a trojan
3	compressnet	Compression Process
5	rje	Remote Job Entry
7	echo	Echo
9	discard	sink null Discard
11	systat	Active Users users Active Users
13	daytime	Daytime (RFC 867)
15	netstat	Network status
17	qotd	Quote of the Day

Figure: Database of well-known TCP and UDP ports.

License

License owner and license key are necessary to run this software. Each customer has a unique license key. To change the license key, click on the "Edit" link. License owner string and license key are not case sensitive both will be checked when you login to the web interface. If license validity time is less than 10 days, a warning window will be displayed.

License data	
Status:	valid license key (edit)
Owner:	TRIAL
License key:	DEMO-0-3252-A202AFA677-4CCF3A88
License type:	unlimited functions, TRIAL version
Valid to:	9 Apr 2005

Figure: License dialog window.

Logout

When you click on the "Logout" menu, the system will try to close your session and free resources.

6 Optimizing and tuning

The heart of the system is MySQL database. This database consumes most of the memory and utilizes the majority of the CPU and disc. For this reason we recommend optimizing the database server. Carefully read MySQL documentation and especially the chapter on "Optimizing the MySQL Server". The MySQL documentation can be obtained from URL http://www.mysql.org/doc/.

On systems with two processors and a 1GB memory we recommend using the following configuration:

```
[mysqld]
skip-locking
#skip-networking
set-variable    = key_buffer=256M
set-variable    = max_allowed_packet=1M
set-variable    = table_cache=1024
set-variable    = sort_buffer=128M
set-variable    = record_buffer=8M
set-variable    = max_allowed_packet=1M
set-variable    = thread_cache=8
set-variable    = myisam_sort_buffer_size=128M
set-variable    = read_buffer=4M
set-variable    = read_rnd_buffer=4M
set-variable    = query_cache_size=32M
set-variable    = query_cache_type=1
set-variable    = bulk_insert_buffer_size=64M
set-variable    = max_heap_table_size=256M
set-variable    = tmp_table_size=256M

#log            = /var/log/mysql/mysql.log
#log-bin        = /var/log/mysql/mysql-bin.log
log-error       = /var/log/mysql/mysql.err
```

In most cases the configuration is in file */etc/mysql/my.cnf*. Don't forget to restart MySQL after making changes.

In case the collector consumes a lot of CPU you can use another server and move several collectors onto the second unit (server).

Appendix 1: Configuring NetFlow Data Export

This appendix is a brief guide to setting up the NetFlow data export (NDE) on a Cisco routers or intelligent L2/L3/L4 switches.If you have problems with the configuration contact your network administrator or Cisco consultant. For devices that run hybrid mode on a Supervisor Engine (Catalyst 65xx series) it is recommended configure IOS NDE on the MSFC card and CatOS NDE on the Supervisor Engine. For more information about setting up netflow, visit http://www.cisco.com/go/netflow.

Configuring NDE on an IOS device

In the configuration mode on the router or MSFC, issue the following to start NetFlow Export.

First enable Cisco Express Forwarding:

```
router(config)# ip cef
router(config)# ip cef distributed
```

And turn on flow accounting for each input interface with the interface command:

```
interface
ip route-cache flow
```

For example:

```
interface FastEthernet0
  ip route-cache flow

interface Serial2/1
  ip route-cache flow
```

It is necessary to enable NetFlow on all interfaces through which traffic (you are interested in) will flow. Now, verify that the router (or switch) is generating flow stats - try command 'show ip cache flow'. Note that for routers with distributed switching (GSR's, 75XX's) the RP cli will only show flows that made it up to the RP. To see flows on the individual linecards use the 'attach' or 'if-con' command and issue the 'sh ip ca fl' on each LC.

Enable the exports of these flows with the global commands:

```
router(config)# ip flow-export version 5
router(config)# ip flow-export destination <ip_address> 2000
router(config)# ip flow-export source FastEthernet0
```

Use the IP address of your NetFlow Collector and configured listening port. UDP port 2000 is used for example.

We recommend using NetFlow version 5, which is the most recent export version supported by Cisco routers. The 'ip flow-export source' command is used to set up the source IP address of the exports sent by the router or switch. NetFlow Collector can filter incoming traffic on this address. If your router uses BGP protocol, you can configure AS to be included in exports with command:

```
router(config)# ip flow-export version 5 [peer-as | origin-as]
```

The following commands break up flows into shorter segments.

```
router(config)# ip flow-cache timeout active 5
router(config)# ip flow-cache timeout inactive 30
```

In enable mode you can see current NetFlow configuration and state.

```
router# show ip flow export
router# show ip cache flow
router# show ip cache verbose flow
```

Configuring NDE on a CatOS device

In privileged mode on the Supervisor Engine enable NDE:

```
switch> (enable) set mls nde <ip_address> 2000
```

Use the IP address of your NetFlow Collector and configured listening port. UDP port 2000 is used as an example.

We recommend using NetFlow version 7, which is the most recent export version supported by Cisco switches.

```
switch> (enable) set mls nde version 7
```

The following command is required to set up flow mask to full flows.

```
switch> (enable) set mls flow full
```

The following commands break up flows into shorter segments.

```
switch> (enable) set mls agingtime long 128
switch> (enable) set mls agingtime 32
```

If you want to account all traffic within the specified VLANs rather then inter VLAN traffic use CatOS 7.2 or higher and issue the following command:

```
switch> (enable) set mls bridged-flow-statistics enable
```

And enable NDE:

```
switch> (enable) set mls nde enable
```

To see current NetFlow configuration and state issue the following commands:

```
switch> (enable) show mls nde
switch> (enable) show mls debug
```

Configuring NDE on a Native IOS device

To configure NDE use the same commands as for the IOS device. In the enable mode on the Supervisor Engine, issue the following, to set up the NetFlow export version.

```
switch(config)# mls nde sender version 7
```

The following commands break up flows into shorter segments.

```
switch(config)# mls aging long 128
switch(config)# mls aging normal 32
```

On the Supervisor Engine 1 issue the following to put full flows into the netfow exports:

```
switch(config)# mls flow ip full
```

If you have a Supervisor Engine 2 or 720 running IOS version 12.1.13(E) or higher, issue the following commands instead:

```
switch(config)# mls flow ip interface-full
switch(config)# mls nde interface
```

Configuring NDE on a 4000 series switch

Configure the switch the same as an IOS device, but instead of command ‘ip route-cache flow’ use command ‘ip route-cache flow infer-fields’. This series requires a Supervisor IV with a NetFlow Services daughter card to support NDE.

Configuring NDE on a Juniper router

Juniper supports flow exports by the routing engine sampling packet headers and aggregating them into flows. Packet sampling is done by defining a firewall filter to accept and sample all traffic, applying that rule to the interface and then configuring the sampling forwarding option.

```
interfaces {
    ge-0/1/0 {
        unit 0 {
            family inet {
                filter {
```

```
                    input all;
                    output all;
                }
                address 192.168.1.1/24;
            }
        }
    }
}

firewall {
    filter all {
        term all {
            then {
                sample;
                accept;
            }
        }
    }
}

forwarding-options {
    sampling {
        input {
            family inet {
                rate 100;
            }
        }
        output {
            cflowd 192.168.1.100 {
                port 2000;
                version 5;
            }
        }
    }
}
```

Other options exist such as aggregated flows, which are detailed at: http://www.juniper.net/.

Appendix 2: Frequently Asked Questions

For clients, who have no knowledge and experience with Linux / Debian

If possible use the self-installing software version, it makes your installation faster, easier and problem free. ISO-CD IMAGE is available at: http://www.caligare.com/netflow/download.php

For clients, who have knowledge and experience with Linux / Debian

Problems encountered with software installation (.deb, .tgz, .rpm) are mostly related to the difference between the Linux environments (library incompatibilities, missing packages, different paths to binaries etc.)

Installation

I can't connect to MySQL database.

When you see this message during installation, type the correct username and password for access to MySQL database. If you are trying to connect to a remote database, check if MySQL is configured for networking (disable option --skip-networking in MySQL configuration file). If you forgot your username and password into MySQL, please, refer to the database documentation (chapter How to Reset the Root Password) or use following steps:

1. Log on to your system as either the Unix root user or as the same user that the mysqld server runs as.
2. Locate the `.pid' file that contains the server's process ID. The exact location and name of this file depends on your distribution, hostname, and configuration. Common locations are `/var/lib/mysql/', `/var/run/mysqld/', and `/usr/local/mysql/data/'. Generally, the filename has the extension of `.pid' and begins with either `mysqld' or your system's hostname. Now you can stop the MySQL server by sending a normal kill (not kill -9) to the mysqld process, using the pathname of the `.pid' file in the following command:

   ```
   kill `cat /mysql-data-directory/host_name.pid`
   ```

   Note the use of back ticks rather than forward quotes with the cat command; these cause the output of cat to be substituted into the kill command.

3. Restart the MySQL server with the special --skip-grant-tables option:

   ```
   mysqld --skip-grant-tables &
   ```

4. Set a new password for the root@localhost MySQL account:

   ```
   mysqladmin -u root flush-privileges password "newpwd"
   ```

Replace ``newpwd'' with the actual root password that you want to use.

5. Restart the MySQL server without any special option.

   ```
   mysqld_safe &
   ```

6. You should now be able to connect using the new password.

Apache configuration file is not found.

If the above message is displayed, you must find and modify the Apache configuration file manually. Configuration filename is mostly httpd.conf and it is stored in the default directory /etc/apache or /etc/apache2. Locate this directory and add to configuration file the following line:

```
Include /etc/netflow/apache.conf
```

In the file "/etc/netflow/apache.conf" there are various options relating to the NetFlow web portion. Don't forget to restart the Apache daemon after modifying its configuration via command:

```
/etc/init.d/apache restart
```

I can't access the web interface.

1. First check if Apache web server is running.

   ```
   ps -ax | grep apache
   ```

2. Check Apache log files.

   ```
   less /var/log/apache/error.log
   ```

   and/or

   ```
   less /var/log/apache/access.log
   ```

3. Check if file /etc/netflow/apache.conf is included in Apache configuration. You can include contents of this file directly into your web server configuration. You can use this file per each virtual host.
4. Check if PHP scripting is enabled in your web server (refer PHP documentation and Apache documentation).

When I tried to restart netflow collector I saw message: "Error: unknown parameter restart"

This message is displayed when you run the command nfcd restart without /etc/init.d/ prefix. Please, run this command with full path. Correct command is:

```
/etc/init.d/nfcd restart
```

Or you can run short nfcd without any parameter, but /etc/init.d/... syntax is preferred. After restarting collector, check your system log file (cat /var/log/syslog).

Web interface

MySQL module isn’t supported by PHP. Check your php.ini file, extensions sections.

This error message is displayed when you haven't installed or activated the MySQL library used by PHP. Try to find the *mysql.so* file by using the following command:

```
find / -name mysql.so
```

When you find this file, activate the extension in your php.ini file (this file is usually located in the directory /etc/php4/apache/) by typing option:

```
extension=mysql.so
```

You can use Midnight command (mc) program to edit this file. If you don't find the *mysql.so* file, try to install a new package *php4-mysql* (package name php4-mysql is used by Debian, in Fedora distribution it will be found with the same or similar to the Debian's name).

Note: PHP must be loaded with MySQL, SNMP and GD extension. Can't open connection into MySQL database; check username, password and MySQL access rights.

This message is displayed when the web part cannot connect into the database (bad username/password or database server hostname not found or database is not running).

1. Check if the php.ini file contains line: *extension=mysql.so*
   - IF YES, please edit file */etc/netflow/nfw.php* and make sure that you have the correct parameters for the database connection (user name, password, database name is nfx)
   - IF NOT, please add line *extension=mysql.so*, save php.ini file and restart your Apache web server.
2. Check if MySQL server is running. In the Linux environment type the following commands:

   ```
   ps ax | grep mysql
   mysql –u root –p
   mysql> quit;
   ```

   If database is not running type the following command:

   ```
   /etc/init.d/mysql start
   ```

3. Based on our PHP knowledge, the PHP module *mysql.so* is probably compiled with an old libmysqlclient version 10. There are several recommendations that might help:

- Try commands:

  ```
  ldconfig
  ldconfig -p | grep mysql
  ```

  Please, send us the output of this command.

- Try restarting Apache.
- Check if your PHP package is the newest version (try upgrading PHP or degrade mysql).
- Send us the output of the following command:

  ```
  rpm –qa
  ```

  (This command will write a list of installed packages on your system – use only for RedHat, SUSE, Fedora distribution).

Can't select MySQL database 'nfx'; check if database exists or you have access rights to use it.

When you ran the nf_install script did you successfully complete step 1? Step 1 creates database and all system tables. Type the following commands to check if step 1 was successfully completed:

```
mysql -u root -p
Password:

mysql> use nfx;
mysql> show tables;
mysql> quit;
```

If software was successfully installed you will see a lot of tables displayed. If it isn't correctly installed then MySQL will write the following information: nfx database doesn't exists. In the Debian installation the password is blank. If you cannot connect into the database due to wrong password you can use the password recovery steps.

When I try to access Data, Trends, I get: Warning: No tables found for selected collector.

Log into the web interface and select menu *Status->Collectors->Detail*. Check if your collector is running (green LED indicator). If you will see a red LED indicator, nfcd process is not running!

If nfcd process is not running, you have to check if your license is OK by going to Help->Licenses. If the License is OK and program is still not running you have to start nfcd process manually. Log into Linux environment and run the following command:

```
/etc/init.d/nfcd restart
```

This command will run the collector(s). You can also see errors or warnings in the system log file (syslog), check if there are any problems with running the collector by using the command:

```
less /var/log/syslog | grep nfc
```

The product is installed and everything seems to be running. However, all the database tables have 0 data in them.

1. Log in into web interface select menu Status->Collectors.
2. Check if your collector is running (green LED indicator). If it is OK, select detail and check all values, you may find there are dropped packets etc.
3. Check if the number of incoming packets is increasing. If not use tcpdump tool, which test receiving NDE packets.

How can I test if netflow collector receives netflow data exports from my Cisco router?

You can use tcpdump tool. Run the following command:

```
tcpdump –n udp
```

You will see all UDP packets that the netflow server receives. You can break tcpdump by typing <Ctrl>+C. If you don't see any packet, check network cable and/or netflow configuration on Cisco router and try debugging netflow exports. If you see incoming packets, but netflow collector still don't receive any packet check your Status->Collector->Detail menu, firewall configuration and system log file (syslog).

Tool tcpdump shows data is coming in. 330 drops where indicated due to bad source IP address in the collector status.

You have to change your device IP address in the menu Options->Devices. The correct IP is IP address from that flows are received. Configure correct source interface on Cisco router or you can use the tcpdump tool for finding correct IP address.

Tool tcpdump shows data is coming in, 150 drops due to bad netflow version in the collector status.

Problem is with unsupported netflow version. Please, configure one of the supported versions on your Cisco router or switch. Supported versions are 1,5,6,7 and 9.

Tool tcpdump shows data is coming in. I did see non-zero DoS state value in the collector status.

In case "DoS state" is non-zero, denial of service protection plug-in blocks data flow, find the source of attack and block it or you can disable this plug-in in the menu "Options->Collectors".

Tool tcpdump shows data is coming in, but 1000 flows indicate corrupted time. Time in exported flows is different then local Linux time.

1. Check if on your Cisco box is valid time via command:

   ```
   show clock
   ```

2. Check if on netflow Linux box is valid time via command:

   ```
   date
   ```

If Cisco and/or Linux time are not synchronized netflow collector drops flows with bad time value. The problem might be in Time Zone set up (information about which time zone you are located in). Please log into Linux environment. In order to set up time zone you have to use the following command:

```
tzsetup -g
```

This command will display recent time zone and ask if you want to change this time zone. If YES, press Y and applications will offer you various continents, cities or countries that you can choose from. (E.g. for United States type in 3, and then type in your time zone). Changes in this setting are saved automatically. When your changes are completed you have to restart your collector using the following command: /etc/init.d/nfcd restart or better, restart your computer via <Ctrl>+<Alt>+<Del>.

To set correct time in the Linux environment you can use date program or you can use the SETUP utility when your computer starts up. If you use date program type the following command:

```
date MMDDhhmmYYYY
```

Where MM is the month number, DD is the day, hh is current hour, mm is current minute and YYYY is the current year. (e.g. *date 030415062005* set up system date is the 4th of March 2005 15:06.)

We recommend use NTP protocol (*ntpdate* utility) instead of manually configured date.

I saw trends results formatted into tables, but no graph is displayed.

You probably haven't installed PHP GD support. Go to the menu “Status->Engine" and check if GDlib is installed.

Can I use more collectors listening on the same port?

Yes, but each collector must have an associated appropriate device. If more collectors share same port to run in one process, can increase CPU utilization; so be careful when using more collectors sharing the same port.

Can I use one collector for more devices?

Yes, but all traffic from these devices will be merged into a common table. It can be useful only for L3/L4 switches, where the L2 switching part exports NetFlow version 7 and the routing engine exports NetFlow version 5. In this case merging these flows into one collector can be very useful; this collector will have complete box traffic.

Is it possible to change the data format for a collector?

No. When you want to change the format, simply delete the collector (all data tables will be dropped!) and re-create it with a new data format or you can disable the old collector and create a new one.

I saw graphs in menu "Data->Graphs", but now they aren't available.

Graphs with type "cache" are removed after 1 day. If you want to save these graphs, select them and click on the "Save cached" button.

When I change the selected table in the trends menu available statistics are changed.

List of available statistics can be changed for different tables, because each collector can define different format of stored data. Check format for each selected collector in the menu "Options->Collectors->Edit".

Other difficulties

If you have any problem with CFI installation or CFI running, please let us know. If you cannot find solution of your problem on this page, please provide us with as many information about your situation and problem as you can.

Detailed information about errors and/or warnings can be found in the system log file (syslog). Please, check if there are any problems, using the following command:

```
less /var/log/syslog | grep nfc
```

Or you can use our debug information collector tool. Run the following command:

```
nf_debug
```

Nf_debug tool send debug information to our support email address. Many companies can have outgoing SMTP traffic blocked and your debug information file can not be sent directly to our email; in this case you have to open the web address: *http://your_netflow_server/netflow/nf_debug.txt* and send us displayed page.

Appendix 3: Network anomalies modules

Network port scanning

The network port scan module detects many suspicious activities as worms, BOTNET scanning attacks, etc. The latest software version detects stations which are scanning the network and looking for network vulnerabilities e.g.: Microsoft WINS, NETBIOS, Microsoft DS, SOCKS, Microsoft SQL, MySQL, web cache, VNC, Microsoft EPMAP and Microsoft terminal services. This module also detects SWIFT, DABBER, QWIN worms and many other unusual activities.

Host port scanning

This network detection module identifies attackers that scan TCP or UDP service ports for vulnerabilities. This module supports only scanning of applications that uses low ports (1-1024).

ICMP flooding

The ICMP flooding detection checks how many ICMP packets the host is sending. If the number of packets exceeds the configured threshold, then the system creates a new anomaly. System recognizes long ICMP messages (>1000B) so that you can configure different thresholds for short ICMP messages and long ICMP messages. Software is capable of detecting unreachable messages (often it signify infection by worm) and other ICMP message types.

TCP/SYN flooding

The TCP/SYN flooding module detects direct or distributed flooding of network with TCP connection requests. This attack is characteristic for distributed denial of service attacks.

Network games detection

The network games detection module uses heuristic methods to detect network games. Many games use the same TCP or UDP port so it is very difficult to say which game was used. The latest version supports the following games: Need for Speed, Diablo, Civilization, Worms 3D, Microsoft DirectX games, Railroad Tycoon, Athena Sword, Unreal, Team Speak, Battlefield 1942, Battle Zone, Age of Empires, Heretic, Hexen, Doom, Call Of Duty, Castle Wolfenstein, Battlefield 2142, MSN Game Zone, Alien vs. Predator, America's Army, Battle.NET, Vietcong, Half-Life and Quake.

Peer to peer application detection

Peer to peer applications waste network bandwidth the most, so detection of these applications is very useful for many administrators, detection of these applications is very, very difficult. Network analysis software uses well-known TCP/UDP ports and some heuristic methods, but in some cases may detect false positives. The latest

version supports detection of the following applications: FastTrack, Kazza, Overnet, Kademlia, Aimster, GNUtella, GNUtella2, WinMX, OpenNapster, Direct Connect, SoulSeek, eDonkey and BitTorrent.

Appendix 4: Third party software components

Our software makes use of several third party libraries, distributed under various licenses.

Apache web server

This product uses software developed by the Apache Software Foundation (http://www.apache.org). This is distributed under the Apache Software License, a copy of which is available at http://www.apache.org/LICENSE.

PHP

This product uses software developed by the PHP Group (http://www.php.net/). This is distributed under the PHP License, a copy of which is available at http://www.php.net/license/3_0.txt.

JPGraph library

This product includes software developed by the Aditus Consulting (http://www.aditus.nu/jpgraph). This is distributed under the JpGraph Professional License, a copy of which is available at http://www.aditus.nu/jpgraph/jpgraph_bulk_license.pdf.